arXiv:1511.08963v2 [math.ST] 14 Dec 2015

LEARNING DIRECTED ACYCLIC GRAPHS WITH PENALIZED NEIGHBOURHOOD REGRESSION

BRYON ARAGAM,[1] ARASH A. AMINI,[2] AND QING ZHOU[3]

University of California, Los Angeles

Abstract. We consider the problem of estimating a directed acyclic graph (DAG) for a multivariate normal distribution from high-dimensional data with $p \gg n$. Our main results establish nonasymptotic deviation bounds on the estimation error, sparsity bounds, and model selection consistency for a penalized least squares estimator under concave regularization. The proofs rely on interpreting the graphical model as a recursive linear structural equation model, which reduces the estimation problem to a series of tractable neighbourhood regressions and allows us to avoid making any assumptions regarding faithfulness. In doing so, we provide some novel techniques for handling general nonidentifiable and nonconvex problems. These techniques are used to guarantee uniform control over a superexponential number of neighbourhood regression problems by exploiting various notions of monotonicity among them. Our results apply to a wide variety of practical situations that allow for arbitrary nondegenerate covariance structures as well as many popular regularizers including the MCP, SCAD, ℓ_0 and ℓ_1.

1. Introduction

Suppose

$$X = (X_1, \dots, X_p) \sim \mathcal{N}_p(0, \Sigma), \tag{1.1}$$

where $X \in \mathbb{R}^p$ and the covariance matrix $\Sigma \in \mathbb{R}^{p \times p}$ is assumed to be positive definite. Then there exist matrices $\widetilde{B}$ and $\widetilde{\Omega}$ such that the following identity holds:

$$X = \widetilde{B}^T X + \widetilde{\varepsilon}, \quad \widetilde{\varepsilon} \sim \mathcal{N}_p(0, \widetilde{\Omega}). \tag{1.2}$$

One can associate a graph to $\widetilde{B}$ by interpreting it as a weighted adjacency matrix. We will say that $\widetilde{B}$ is, or represents, a DAG if this graph is directed and acyclic. In the Gaussian setting, it is always possible to find a representation as in (1.2) where $\widetilde{B}$ is a DAG and $\widetilde{\Omega}$ is diagonal. This decomposition is often referred to as a linear *structural equation model* (SEM) for X, and extends well beyond the Gaussian setting considered here.

E-mail address: [1]bryon@stat.ucla.edu, [2]aaamini@stat.ucla.edu, [3]zhou@stat.ucla.edu.

A fundamental property of this model is that the parameters $(\widetilde{B}, \widetilde{\Omega})$ are statistically nonidentifiable. In fact, it is easy to construct a DAG satisfying (1.2): For each $j < p$, project X_j onto the linear span of $\{X_{j+1}, \ldots, X_p\}$. However, if we re-order the variables and re-apply this procedure, we may end up with a completely different set of structural equations. Thus, without additional information this model is nonidentifiable with up to $p!$ distinct values for $(\widetilde{B}, \widetilde{\Omega})$, one for each permutation of the variables.

Given an ordering on the nodes X_j, the SEM (1.2) factorizes the joint density of X into a product of conditional densities. Since any joint density admits a factorization which can be represented by a DAG, structure learning of DAGs from observed data is a very general and fundamental problem in statistics. Even for the Gaussian case considered in this work, however, the problem of DAG estimation from high-dimensional data remains very challenging. This is largely due to the complexity of the problem: In addition to the nonidentifiability, optimization over the space of DAGs is well-known to be nonconvex and NP-hard in general. This leads to some important questions:

(1) What can we say about the class of DAGs satisfying (1.2)?
(2) Under what conditions can we estimate *any* of these DAGs?
(3) Assuming we can do this, what kind of guarantees can we make about the sparsity of the estimated DAG?

In this paper, we attempt to address these questions, with a particular focus on providing support recovery and sparsity guarantees.

1.1. **Overview.** Given an $n \times p$ random matrix $\mathbf{X}$ whose rows are i.i.d. drawn according to $\mathcal{N}_p(0, \Sigma)$, define an objective function by

$$Q(B) = \frac{1}{2n} \|\mathbf{X} - \mathbf{X}B\|_2^2 + \rho_\lambda(B), \tag{1.3}$$

where $\|\cdot\|_2$ denotes the matrix Frobenius norm and ρ_λ belongs to a suitably chosen class of regularizers, parametrized by the *regularization parameter* λ. Letting

$$\mathbb{D}_p = \{B \in \mathbb{R}^{p \times p} : B \text{ is a DAG}\}, \tag{1.4}$$

we study the following estimator:

$$\widehat{B} \in \operatorname*{arg\,min}_{B \in \mathbb{D}_p} Q(B). \tag{1.5}$$

Our main result is to establish the consistency of $\widehat{B}$ in both model selection and parameter estimation in a high-dimensional setting, under no assumptions on the ordering, as well as to identify which of the many DAGs satisfying (1.2) $\widehat{B}$ approximates. Our results allow for arbitrary dependence between the components of X—including the case where the variances of $\widetilde{\varepsilon}$ are not the same—as well as a general class of regularizers, including the ℓ_1 and ℓ_0 penalties, MCP (Zhang, 2010), and SCAD (Fan and Li, 2001).

This problem must be contrasted with the comparatively simpler problem of estimating an undirected Gaussian graphical model, which has received much more attention. Indeed, Meinshausen and Bühlmann (2006) showed how to reduce the estimation problem for the Gaussian graphical model to a sequence of so-called *neighbourhood regressions*. The procedure is conceptually simple: For each j, regress X_j onto the rest of the variables X_{-j}. This procedure works for other classes of undirected graphical models, including Ising models (Ravikumar et al., 2010) and discrete graphical models (Loh and Wainwright, 2013). As we will show, this intuition also works for *directed* models with one important caveat: Instead of p neighbourhoods, one needs to control a superexponential number of neighbourhoods, which requires new techniques and subtler arguments.

In choosing the loss function (1.3), we could have used the negative log-likelihood of the Gaussian model (1.1), as in van de Geer and Bühlmann (2013), in place of the least-squares (LS) term $\frac{1}{2n}\|\mathbf{X} - \mathbf{X}B\|_2^2$. We have chosen LS due to its simplicity and its nice interpretation in terms of neighbourhood regressions, allowing us to outline our main arguments in the more familiar context of linear regression. However, what is necessary on the loss is that (1.3) factors into p neighborhood problems (for each fixed ordering of the variables), as detailed in Lemma 5.1, a property that holds for other losses as well. In addition, the LS loss has the potential for generalization to the subgaussian setting where there is no specific likelihood. The convexity of the LS loss also allows for fast and accurate algorithms in high-dimensions: In particular, one can use coordinate descent in order to approximate $\widehat{B}$ efficiently, even when p is in the thousands and $p > n$. As the focus of this work is mainly theoretical, we omit further discussion of such algorithmic concerns. For a more thorough treatment of computational aspects of DAG estimation, see Aragam and Zhou (2015).

The organization of the rest of this paper is as follows: In the rest of this section, we review previous work and outline our contributions at a high-level. In Section 2, we give some intuition behind the estimator (1.5) and cover some necessary background material. In Section 3 we establish nonasymptotic bounds on the estimation error and sparsity of $\widehat{B}$. Then, in Section 4, we strengthen these results by providing explicit guarantees for support recovery. Finally, we provide proofs of the main results in Section 5, with the proofs of various technical results deferred to the Appendix.

1.2. **Previous work.** The interpretation of directed graphical models in terms of linear structural equations has been shown to lead to fast, efficient algorithms for penalized estimation of DAGs that extend even beyond the Gaussian set-up we consider presently. Some early contributions include Schmidt et al. (2007) and Shojaie and Michailidis (2010). More recently, Fu and Zhou (2013) introduced a novel method for inferring causal relationships based on experimental data, which was later extended to discrete data in Fu

et al. (2014) and also to structure learning of Gaussian Bayesian networks from observational data in Aragam and Zhou (2015).

Compared to these methodological developments, theoretical progress in high-dimensions has been slower. van de Geer and Bühlmann (2013) provide the first high-dimensional results, establishing bounds on the ℓ_2-error and the number of edges in $\widehat{B}$ for a thresholded ℓ_0-penalized maximum likelihood estimator. They also consider the equivariant case, which is identifiable (Peters and Bühlmann, 2013), and under some conditions are able to guarantee recovery of the equivariant permutation. In contrast, we focus on the fully nonidentifiable case and show how equivariance is a special case of the framework considered in the present work. Furthermore, our work covers a wide spectrum of concave regularization including both ℓ_0 and ℓ_1 as boundary cases. A more detailed comparison between the two approaches is provided in Section 3.3.4.

There has also been progress using multi-stage methods that separate the learning procedure into several decoupled steps. Loh and Bühlmann (2014) show how to use the so-called moral graph to restrict the search space, and then use a re-weighted least-squares estimator to reconstruct a linear SEM. They also provide some interesting results regarding identifiability of such models. Bühlmann et al. (2014) propose a three-step method for learning nonlinear SEM that uses preliminary neighbourhood search, order search, and finally sparse additive regression. In contrast to these results, which make strong use of various assumptions that guarantee model identifiability, the proof technique used in the present work is explicitly designed to overcome the nonidentifiability of the general Gaussian model. Furthermore, our method consists of a single learning step that effectively combines neighbourhood search, order search, and regression by minimizing one penalized loss function.

Notation and terminology. Arbitrary constants which may not be the same from line to line will be denoted by c_1, c_2, etc. The standard ℓ_q norms for $q \in [0, \infty]$ will be denoted by $\|\cdot\|_q$. For a matrix $A \in \mathbb{R}^{n\times p}$, $\|A\|_q$ is its ℓ_q norm, viewing A as a vector in $\mathbb{R}^{np}$, and so, in particular, the Frobenius norm on matrices is denoted by $\|\cdot\|_2$. The maximum and minimum eigenvalues of a matrix A are denoted by $r_{\max}(A)$ and $r_{\min}(A)$, respectively. The support of $A = (a_{ij})$ is defined by $\operatorname{supp}(A) := \{(i,j) : a_{ij} \neq 0\}$ and the cardinality of a set by $|\cdot|$. Given a permutation π, P_π denotes the associated permutation operator on matrices: For any matrix A, $P_\pi A$ is the matrix obtained by permuting the rows and columns of A according to π, so that $(P_\pi A)_{ij} = a_{\pi(i)\pi(j)}$. For any integer m, we define $[m] = \{1, \dots, m\}$ and $[m]_j = [m] - \{j\}$. For a vector $v \in \mathbb{R}^m$ and a subset $S \subset [m]$, we let $v_S \in \mathbb{R}^{|S|}$ denote the restriction of v to the components in S. For a matrix $A \in \mathbb{R}^{n\times m}$, $A_S \in \mathbb{R}^{n\times|S|}$ denotes *column-wise* restriction to the columns in S. Given two quantities X and Y, which may depend on n and Σ, we will write $X \lesssim Y$ to mean there exists a constant $a > 0$—independent

of n and Σ—such that $X \leq aY$, and analogously for $X \gtrsim Y$. A positive definite matrix is denoted as $\Sigma \succ 0$.

We write $\mathbf{X} \overset{\text{iid}}{\sim} \mathcal{N}_p(0, \Sigma)$ for a random matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ to mean that the *rows* are i.i.d. draws from $\mathcal{N}_p(0, \Sigma)$. As a general rule, boldface type is reserved for *random* quantities that also depend on the sample size n; for example a random matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ or the response $\mathbf{y} \in \mathbb{R}^n$ in a linear model. Events defined on some probability space will usually be denoted by script font, e.g. $\mathcal{E}$, $\mathcal{F}$, etc.

1.3. **Nonidentifiability and permutation equivalence.** We start with an example to highlight some key ideas.

Example 1.1. Consider the following covariance matrix:

$$\Sigma = \begin{pmatrix} 6 & 4 & -6 & -30 \\ 4 & 4 & -4 & -20 \\ -6 & -4 & 7 & 39 \\ -30 & -20 & 39 & 234 \end{pmatrix}. \tag{1.6}$$

Given Σ, we can construct matrices $(\widetilde{B}, \widetilde{\Omega})$ that satisfy (1.2) by the following procedure:

Project X_2 onto nothing,
Project X_1 onto X_2,
Project X_3 onto X_2 and X_1,
Project X_4 onto X_2, X_1, and X_3,

This induces an ordering $\prec$ on the variables given by $X_4 \prec X_3 \prec X_1 \prec X_2$, wherein each X_j is projected onto the nodes after it under $\prec$. This ordering induces a permutation π_1 defined by

$$\pi_1(1) = 4, \quad \pi_1(2) = 3, \quad \pi_1(3) = 1, \quad \pi_1(4) = 2.$$

The coefficients of each linear projection and the resulting residual variances lead to a set of structural equations as in (1.2).

For example, denoting the parameters obtained in this way by $\widetilde{B} = \widetilde{B}(\pi_1)$ and $\widetilde{\Omega} = \widetilde{\Omega}(\pi_1)$, we obtain the following parameters for (1.6):

$$\widetilde{B}(\pi_1) = \begin{pmatrix} 0 & 0 & -1 & 4 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 9 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \widetilde{\Omega}(\pi_1) = \begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 4 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 3 \end{pmatrix}. \tag{1.7}$$

This DAG is depicted on the left in Figure 1.

If we choose a different ordering, say $X_4 \prec X_1 \prec X_2 \prec X_3$, we obtain a different set of parameters:

$$\widetilde{B}(\pi_2) = \begin{pmatrix} 0 & 0 & 0 & 4 \\ \frac{1}{3} & 0 & 0 & 0 \\ -\frac{2}{3} & -\frac{4}{7} & 0 & 9 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \widetilde{\Omega}(\pi_2) = \begin{pmatrix} \frac{2}{3} & 0 & 0 & 0 \\ 0 & \frac{12}{7} & 0 & 0 \\ 0 & 0 & 7 & 0 \\ 0 & 0 & 0 & 3 \end{pmatrix}, \tag{1.8}$$

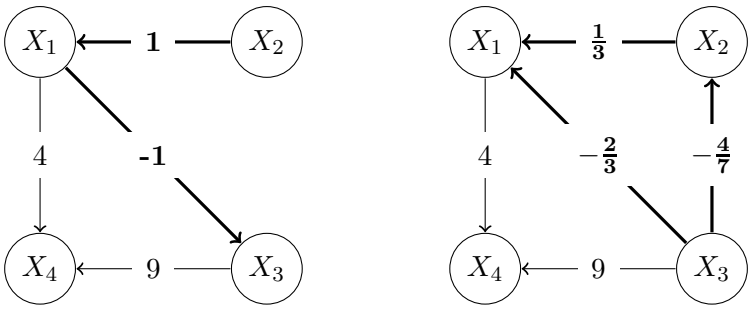

FIGURE 1. Two equivalent DAGs with different edge weights, as indicated by the edges highlighted in bold.

where π_2 is the permutation induced by this ordering. This DAG is depicted on the right in Figure 1. The fact that the incoming edges for X_4 are the same in both cases is consistent with the observation that under both orderings, we are projecting X_4 onto $\{X_1, X_2, X_3\}$.

In the previous example, both (1.7) and (1.8) lead to a linear model as in (1.2) with $\text{cov}(X) = \Sigma$, but with distinct numbers of edges. Two more key points from Example 1.1 are worth emphasizing:

- Both DAGs $\widetilde{B}(\pi_j)$ $(j = 1, 2)$ are very sparse (i.e. as measured by the number of nonzero entries in $\widetilde{B}(\pi_j)$) even though Σ itself is not sparse at all,
- The conditional variances given by $\widetilde{\Omega}(\pi_j)$ $(j = 1, 2)$ are different in general, and in fact in this example there does not exist a DAG $\widetilde{B}$ satisfying (1.2) that has the same variances for each node.

These observations motivate much of the developments in the sequel, and lead us to consider the following definition:

Definition 1.1. Let $\mathbb{R}^p_+$ denote the space of $p \times p$ diagonal matrices with positive entries on the diagonal. Given a covariance matrix $\Sigma \succ 0$ and $X \sim \mathcal{N}_p(0, \Sigma)$, define the *equivalence class* of Σ to be

$$\mathfrak{D}(\Sigma) := \{\widetilde{B} \in \mathbb{D}_p : \widetilde{B} \text{ satisfies (1.2) for some } \widetilde{\Omega} \in \mathbb{R}^p_+\},$$

where $\mathbb{D}_p$ is the space of DAGs as defined in (1.4). The DAGs in $\mathfrak{D}(\Sigma)$ will be called *equivalent*.

We will now show that this equivalence class can be constructed explicitly. Let $\mathbb{S}_p$ denote the class of permutations on p elements and $\pi \in \mathbb{S}_p$ be a fixed permutation. Write $\Gamma := \Sigma^{-1}$. We may use the (modified) Cholesky decomposition to write $P_\pi \Gamma$ uniquely as $P_\pi \Gamma = (I - L)D^{-1}(I - L)^T$ where L is strictly lower triangular and $D \in \mathbb{R}^p_+$. Define

$$\widetilde{B}(\pi) := P_{\pi^{-1}} L, \quad \widetilde{\Omega}(\pi) := P_{\pi^{-1}} D. \tag{1.9}$$

Since an adjacency matrix represents a DAG if and only if it is permutation-similar to a strictly lower triangular matrix, $\widetilde{B}(\pi)$ is indeed a DAG. We refer to $\widetilde{B}(\pi)$ as the DAG associated with π.

Lemma 1.1. *Suppose* $X \sim \mathcal{N}_p(0, \Sigma)$ *with* $\Sigma \succ 0$. *Then* $\mathfrak{D}(\Sigma) = \{\widetilde{B}(\pi) : \pi \in \mathbb{S}_p\}$, *where* $\widetilde{B}(\pi)$ *is defined by* (1.9).

This characterization of the equivalence class will play an important role in the rest of the paper. Therefore, without further qualification, we shall always write an arbitrary element of $\mathfrak{D}(\Sigma)$ as $\widetilde{B}(\pi)$. The columns of $\widetilde{B}(\pi)$ will be denoted by $\widetilde{\beta}_j(\pi)$, and the jth diagonal element of $\widetilde{\Omega}(\pi)$ will be denoted by $\widetilde{\omega}_j^2(\pi)$. It follows from these definitions and (1.2) that

$$X_j = \widetilde{\beta}_j(\pi)^T X + \widetilde{\varepsilon}_j(\pi), \quad \text{where} \quad \widetilde{\varepsilon}_j(\pi) \sim \mathcal{N}(0, \widetilde{\omega}_j^2(\pi)),$$

for $j = 1, \ldots, p$.

1.4. **Contributions.** Recall the definitions of $\widehat{B}$ in (1.5) and $\widetilde{B}(\pi)$ in (1.9). Our goal is (a) To show that there exists a (random) permutation $\widehat{\pi}$ such that $\widehat{B}$ is close to $\widetilde{B}(\widehat{\pi})$, and (b) To compare the sparsity of $\widehat{B}$ and $\widetilde{B}(\widehat{\pi})$ to some *fixed* DAG $\widetilde{B}(\pi_0)$. In the following $\kappa > 0$ is a constant that depends on the eigenvalues of Σ, and could be different in each occurrence. Broadly speaking, our results can be summarized as follows:

(1) **Deviation bounds:** We show that there exists a random permutation $\widehat{\pi}$ such that the deviation $\widehat{B} - \widetilde{B}(\widehat{\pi})$ is small in the ℓ_2 and ℓ_1 sense. More specifically, we show that for properly chosen values of the regularization parameter λ, we have

$$\|\widehat{B} - \widetilde{B}(\widehat{\pi})\|_r \le \kappa\, \rho'_\lambda(0+) \|\widetilde{B}(\widehat{\pi})\|_0^{1/r}, \quad r = 1, 2, \tag{1.10}$$

where $\rho'_\lambda(0+)$ is the right derivative of ρ_λ at zero. Here, $\widehat{\pi}$ is a permutation such that $P_{\widehat{\pi}}\widehat{B}$ is lower triangular, which can be obtained immediately once $\widehat{B}$ is found. Moreover, using standard bounds (see, e.g., Bickel et al. (2009)), these can be readily extended to ℓ_r bounds for $1 \le r \le 2$.

(2) **Sparsity bounds:** The deviation bound (1.10) shows that $\widehat{B}$ is close to $\widetilde{B}(\widehat{\pi})$, which depends on the random permutation $\widehat{\pi}$. One could ask how sparse $\widetilde{B}(\widehat{\pi})$ is. We answer this question by comparing the sparsity of $\widetilde{B}(\widehat{\pi})$ to a *minimum-trace* DAG $\widetilde{B}(\pi_0)$, defined via a permutation π_0 that minimizes $\pi \mapsto \operatorname{tr} \widetilde{\Omega}(\pi)$. More specifically, we show that with high probability,

$$\|\widetilde{B}(\widehat{\pi})\|_0 \le \kappa\, \|\widetilde{B}(\pi_0)\|_0, \quad \forall \pi_0 \in \arg\min_{\pi}\ \operatorname{tr} \widetilde{\Omega}(\pi). \tag{1.11}$$

In other words, the sparsity of the DAG associated with the random permutation $\widehat{\pi}$ is within a constant factor of the sparsity of a minimum-trace DAG. This observation is novel and points to some interesting new research directions for studying directed acyclic graphs.

(3) **Model selection consistency:** We prove the model selection consistency of $\widehat{B}$, by showing that the event $\{\operatorname{supp}(\widehat{B}) \neq \operatorname{supp}(\widetilde{B}(\widehat{\pi}))\}$ has vanishingly small probability. Combined with (1.11), this leads to the same guarantee for the sparsity of $\widehat{B}$ in terms of the sparsity of the minimum-trace DAG.

All of our results will be nonasymptotic with an emphasis on the case $p \gg n$. The bound on model selection failure, along with the deviation bounds (1.10), will turn out to be special cases of a general series of results regarding uniform consistency of a collection of neighbourhood regression problems, to be described in more detail in Sections 4 and 5. Compared to the existing literature, the main contributions of this paper are (a) Model selection guarantees when $p \gg n$, (b) ℓ_r bounds on the estimation error for $1 \leq r \leq 2$, (c) Allowing both of these for a very general class of regularizers, including the ℓ_1 and ℓ_0 penalties, MCP, and SCAD. Finally, the proof technique via model selection exponents and minimum-trace DAGs is a novel contribution to the literature on graphical models, and applies to settings beyond those considered herein.

2. Background and preliminaries

In this section we provide some intuition behind the estimator (1.5) and establish some assumptions on Σ and the regularizer ρ_λ.

2.1. Global and restricted minimizers.

Recall that $\mathbb{D}_p$ is the space of $p \times p$ real matrices that represent DAGs when interpreted as weighted adjacency matrices. For each permutation $\pi \in \mathbb{S}_p$, define a subset of $\mathbb{D}_p$ by

$$\mathbb{D}_p[\pi] = \{B \in \mathbb{D}_p : P_\pi B \text{ is lower triangular}\}.$$

For any fixed permutation π, the structure of the DAGs in $\mathbb{D}_p[\pi]$ can be described as follows: A DAG $B = [\,\beta_1 \,|\, \cdots \,|\, \beta_p\,] \in \mathbb{D}_p$ is in $\mathbb{D}_p[\pi]$ if and only if for each $j = 1, \ldots, p$ it holds that $\operatorname{supp}(\beta_j) \subset S_j(\pi)$, where

$$S_j(\pi) := \{k : \pi^{-1}(k) > \pi^{-1}(j)\} \tag{2.1}$$

consists of the nodes X_k that come after X_j under the ordering $X_{\pi(i)} \prec X_{\pi(i+1)}$ for $i = 1, \ldots, p-1$. In other words, for each node X_j, the permutation π defines a unique set of candidate parents given by (2.1), and if $B \in \mathbb{D}_p[\pi]$, then for each j the parent set of β_j must come from $S_j(\pi)$. Moreover, note that $\mathbb{D}_p = \bigcup_\pi \mathbb{D}_p[\pi]$ and $\widetilde{B}(\pi) \in \mathbb{D}_p[\pi]$ for every π. It follows that for each $\pi \in \mathbb{S}_p$, $\operatorname{supp}(\widetilde{\beta}_j(\pi)) \subset S_j(\pi)$ for each $j = 1, \ldots, p$.

Recall the objective function defined in (1.3) and the associated estimator (1.5). The spaces $\mathbb{D}_p[\pi]$ help give us some intuition into what is going on with the estimator $\widehat{B}$: Suppose we had oracle knowledge of some "optimal" or "best" permutation π_0 in advance. For example, π_0 might be a permutation that minimizes the number of edges amongst all of the equivalent DAGs. Then we can ignore any DAG that is not consistent with π_0,

which is tantamount to restricting our search space to $\mathbb{D}_p[\pi_0]$. Thus, estimating the "oracle DAG" $\widetilde{B}(\pi_0)$ reduces to a simple autoregressive model, which is equivalent to minimizing $Q(B)$ with the additional constraint that $B \in \mathbb{D}_p[\pi_0]$. This motivates the following general definition:

Definition 2.1. A *restricted global minimizer* of Q is

$$\widehat{B}(\pi) \in \operatorname*{arg\,min}_{B \in \mathbb{D}_p[\pi]} Q(B). \tag{2.2}$$

The columns of $\widehat{B}(\pi)$ will be denoted by $\widehat{\beta}_j(\pi)$, $j = 1, \ldots, p$.

Obviously, $\widehat{B}$ must also be a restricted global minimizer for some π. Unfortunately, we do not know in advance which permutation(s) this corresponds to, i.e. for which π we have $\widehat{B} = \widehat{B}(\pi)$. This leads us to consider the collection of all such permutations, formalized by the following definition.

Definition 2.2. The *collection of estimated permutations* is

$$\widehat{\mathbb{S}}_p := \operatorname*{arg\,min}_{\pi \in \mathbb{S}_p} Q(\widehat{B}(\pi)).$$

An arbitrary element of $\widehat{\mathbb{S}}_p$ will be denoted by $\widehat{\pi}$.

If $\widehat{\pi} \in \widehat{\mathbb{S}}_p$, then the following statements are equivalent: (i) $P_{\widehat{\pi}}\widehat{B}$ is lower triangular, (ii) $\widehat{B} \in \mathbb{D}_p[\widehat{\pi}]$, and (iii) $\widehat{B} = \widehat{B}(\widehat{\pi})$. Since $\widehat{B}$ depends on the random matrix $\mathbf{X}$, the permutation $\widehat{\pi}$ is also a random quantity. This randomness will play a pivotal role in complicating the arguments that appear in the sequel. In practical applications, once a solution $\widehat{B}$ to (1.5) is found, $\widehat{\mathbb{S}}_p$ can be obtained as a byproduct since any DAG must be compatible with at least one ordering among its nodes.

2.2. **Assumptions on Σ.** Recall the equivalence class $\mathfrak{D}(\Sigma)$ defined in Definition 1.1. Our only assumption on Σ ensures that $\mathfrak{D}(\Sigma)$ is well-defined.

Condition 2.1. Σ is positive definite, i.e., $r_{\min}(\Sigma) > 0$.

Define

$$\sigma^2_{\max} := \max_{1 \le j \le p} \operatorname{var}(X_j) \le r_{\max}(\Sigma). \tag{2.3}$$

In the sequel, we will use the symbol κ to indicate a constant that depends on Σ: For example, $\kappa_1 = \kappa_1(\Sigma)$, etc.

Our results will depend on two quantities that are familiar from the regression literature: The sparsity level and the signal strength. The difference now is that instead of quantifying these for a single, "true" parameter, we need to consider the entire equivalence class $\mathfrak{D}(\Sigma)$.

Recall that we denote the jth column of $\widetilde{B}(\pi)$ by $\widetilde{\beta}_j(\pi)$. For any $u \in \mathbb{R}^p$, let $\tau_*(u) := \min\{|u_j| : u_j \neq 0\}$.

Definition 2.3. For any Σ, let

$$d = d(\Sigma) := \sup_{1\le j\le p} \sup_{\pi\in\mathbb{S}_p} \|\widetilde{\beta}_j(\pi)\|_0, \tag{2.4}$$

$$\tau_* = \tau_*(\Sigma) := \inf_{1\le j\le p} \inf_{\pi\in\mathbb{S}_p} \tau_*(\widetilde{\beta}_j(\pi)). \tag{2.5}$$

In other words, d is the size of the largest parent set within any DAG in $\mathfrak{D}(\Sigma)$, and $\tau_*(\Sigma) \in (0,\infty)$ is the smallest nonzero coefficient in absolute value of any DAG in $\mathfrak{D}(\Sigma)$. The parameter d measures the relative sparsity of Σ through its equivalence class, and τ_* measures the minimum signal strength in terms of its SEM coefficients. Both d and τ_* are allowed to depend on n, and in particular d may diverge along with n.

2.3. **Choice of regularizer.** We consider *coordinate-separable* regularizers, which means that we can express $\rho_\lambda(B)$, with some abuse of notation, as

$$\rho_\lambda(B) = \sum_{i,j} \rho_\lambda(|\beta_{ij}|)$$

where $\rho_\lambda : [0,\infty) \to [0,\infty)$ is a univariate regularizer. A key feature of our analysis is to provide some insight into how ρ_λ affects the solution to (1.5), so we will not specify a particular regularizer in advance. Instead, we make the following minimal assumptions on ρ_λ:

Condition 2.2. The regularizer ρ_λ is concave, nondecreasing, right-differentiable at zero with $0 < \rho'_\lambda(0+) < \infty$ and satisfies $\rho_\lambda(0) = 0$.

Condition 2.3. There exist constants $\underline{\rho}_0 \ge 0$ and $\underline{\rho}_1 \ge 0$, independent of λ, such that

$$\rho_\lambda(x) \ge \min\{\underline{\rho}_1 \lambda x, \underline{\rho}_0 \lambda^2\}. \tag{2.6}$$

As an elementary consequence of Condition 2.2, note that ρ_λ must also be subadditive. Condition 2.3 says that ρ_λ can be bounded below by a capped-ℓ_1 penalty: It is always true that a concave function can be bounded below by a capped-ℓ_1 penalty, and Condition 2.3 simply normalizes this capped-ℓ_1 penalty in terms of λ. These conditions allow for most concave penalties, including the MCP and SCAD, along with the familiar ℓ_1 penalty.

Finally, we will also be interested in penalties that "approximate" the ℓ_0 penalty in a sense that is made precise by the following definition:

Definition 2.4. A regularizer ρ_λ is called ℓ_0-compatible if there exists a constant $\overline{\rho}_0 \ge 0$, independent of λ, such that $\rho_\lambda(x) \le \overline{\rho}_0 \lambda^2$ for all $x \ge 0$.

When ρ_λ is ℓ_0-compatible, we can bound $\rho_\lambda(B)$ by the ℓ_0 penalty, i.e. $\rho_\lambda(B) \le \overline{\rho}_0 \lambda^2 \|B\|_0$ for any $B \in \mathbb{D}_p$.

Example 2.1 (MCP). The *minimax concave penalty*, introduced in Zhang (2010), is given by

$$\rho_\lambda(x;\gamma) := \lambda\left(x - \frac{x^2}{2\lambda\gamma}\right) 1(x < \lambda\gamma) + \frac{\lambda^2\gamma}{2} 1(x \ge \lambda\gamma). \tag{2.7}$$

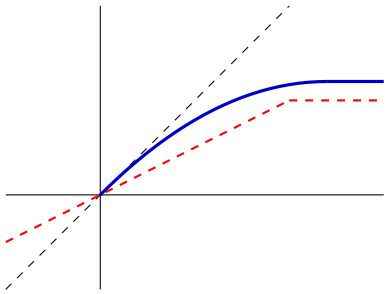

FIGURE 2. Illustration of Conditions 2.2 and 2.3 for MCP with $\lambda = 1$ and $\gamma = 1.2$. The dashed red line highlights one possible lower bound on ρ_λ via a capped-ℓ_1 penalty (Condition 2.3) and the dashed black line illustrates the upper bound $\rho_\lambda(x) \leq \rho'_\lambda(0+)x$ that is implied by Condition 2.2.

The MCP satisfies Conditions 2.2 and 2.3 with $\rho'_\lambda(0+) = \lambda$, $\underline{\rho}_1 = 1/2$, and $\underline{\rho}_0 = \gamma/2$. Furthermore, it is ℓ_0-compatible with $\overline{\rho}_0 = \gamma/2$. See Figure 2.

Example 2.2 (ℓ_1 penalty). The ℓ_1 penalty, $\rho_\lambda(x) = \lambda x$, also satisfies Conditions 2.2 and 2.3 with $\rho'_\lambda(0+) = \lambda$, $\underline{\rho}_1 = 1$, and $\underline{\rho}_0 \in [0, \infty)$. The ℓ_1 penalty is not, however, ℓ_0-compatible.

Example 2.3 (ℓ_0 penalty). The ℓ_0 penalty, $\rho_\lambda(x) = (\lambda^2/2)1(x \neq 0)$, satisfies Condition 2.3 with $\underline{\rho}_1 \in [0, \infty)$ and $\underline{\rho}_0 = 1/2$, but it fails to satisfy Condition 2.2 due to the discontinuity at the origin. Of course, the ℓ_0 penalty is trivially ℓ_0-compatible with $\overline{\rho}_0 = 1/2$.

Remark 2.1. Even though the ℓ_0 penalty does not satisfy Condition 2.2, all the results in this paper still apply to the ℓ_0 penalty under possibly different values of the constants involved. This requires a small modification to the overall argument, which is discussed in Section 6.1.

3. Deviation bounds and sparsity

Our main results are divided into three separate theorems: (i) Bounds on the estimation error $\|\widehat{B} - \widetilde{B}(\widehat{\pi})\|_r$, (ii) Bounds on the sparsity of $\widehat{B}$ and $\widetilde{B}(\widehat{\pi})$, and (iii) Bounds on the probability of false selection. The former results can be applied in a very general setting under fairly weak assumptions, whereas the latter result for model selection requires a more abstract set-up with (possibly) stronger assumptions. The present section addresses the details of (i) and (ii), while (iii) is deferred until the next section.

3.1. **Deviation bounds.**

Our first result bounds the estimation errors for each column, $\|\widehat{\beta}_j - \widetilde{\beta}_j(\widehat{\pi})\|_r$:

Theorem 3.1. *Suppose Conditions 2.1, 2.2, and 2.3 hold. Then there exist positive constants $\kappa_1 = \kappa_1(\Sigma)$ and $\kappa_2 = \kappa_2(\Sigma)$ such that if*

$$n > \kappa_1 \, d \log p, \quad \lambda \geq \kappa_2 \sqrt{\frac{(d+1)\log p}{n}}$$

then for $r = 1, 2$,

$$\|\widehat{B} - \widetilde{B}(\widehat{\pi})\|_r \lesssim \frac{\rho'_\lambda(0+)}{r_{\min}(\Sigma)} \|\widetilde{B}(\widehat{\pi})\|_0^{1/r}, \text{ and} \tag{3.1}$$

$$\|\widehat{\beta}_j - \widetilde{\beta}_j(\widehat{\pi})\|_r \lesssim \frac{\rho'_\lambda(0+)}{r_{\min}(\Sigma)} \|\widetilde{\beta}_j(\widehat{\pi})\|_0^{1/r} \quad \forall\, 1 \leq j \leq p, \tag{3.2}$$

with probability at least $1 - c_1 \exp(-c_2(d+1)\log p)$.

In fact, something much stronger is true: The same ℓ_r bounds in Theorem 3.1 apply uniformly to the deviations $\widehat{B}(\pi) - \widetilde{B}(\pi)$ and $\widehat{\beta}_j(\pi) - \widetilde{\beta}_j(\pi)$ for all $\pi \in \mathbb{S}_p$ (see Proposition B.1). Moreover, using standard interpolation bounds for ℓ_r norms, this result can easily be extended to bounds for all $1 \leq r \leq 2$. See, for example, the proof of Theorem 7.1 in Bickel et al. (2009).

3.2. **Sparsity.** It is worth emphasizing that Theorem 3.1 holds whenever Σ is positive definite, and that no additional assumptions on the population-level parameters are needed. In particular, we do not need any kind of "beta-min" condition on τ_*. Under stronger conditions, we can say more about the DAG $\widetilde{B}(\widehat{\pi})$. In this section we show how we can control the number of edges in $\widetilde{B}(\widehat{\pi})$ under some assumptions on the *signal strength* and a *minimum-trace DAG*.

Condition 3.1 (Signal strength). The minimum signal τ_* satisfies

$$\rho_\lambda(\tau_*) \geq a_1 \frac{\rho'_\lambda(0+)^2}{r_{\min}(\Sigma)} \quad \text{for some } a_1 > 2.$$

This condition is discussed in more detail in Section 3.3.1. Our final condition involves a so-called *minimum-trace DAG*:

Definition 3.1. A *minimum-trace permutation* is any

$$\pi_0 \in \operatorname*{arg\,min}_{\pi \in \mathbb{S}_p} \operatorname{tr} \widetilde{\Omega}(\pi).$$

The corresponding DAG $\widetilde{B}(\pi_0)$ is called a *minimum-trace DAG*.

Condition 3.2 (Minimum-trace). There exists a minimum-trace permutation π_0 such that

$$\frac{\rho_\lambda(\widetilde{B}(\pi_0))}{\operatorname{tr} \widetilde{\Omega}(\pi_0)} \geq a_2 \sqrt{\frac{(d+1)\log p}{n}} \quad \text{for some } a_2 > 0.$$

In light of the lower bound on λ required by Theorem 3.1, Condition 3.2 can be rewritten as $\rho_\lambda(\widetilde{B}(\pi_0)) \gtrsim (\lambda/\kappa_2) \cdot \operatorname{tr} \widetilde{\Omega}(\pi_0)$ which puts a lower bound on the penalty, as measured by the minimum-trace DAG $\widetilde{B}(\pi_0)$.

The following result shows how the sparsity of $\widehat{B}$ is essentially sandwiched between that of $\widetilde{B}(\widehat{\pi})$ and $\widetilde{B}(\pi_0)$:

Theorem 3.2. *Assume the conditions of Theorem 3.1 hold along with Conditions 3.1 and 3.2. Then*

$$\rho_\lambda(\widetilde{B}(\widehat{\pi})) \lesssim \rho_\lambda(\widehat{B}) \lesssim \rho_\lambda(\widetilde{B}(\pi_0)) \tag{3.3}$$

with probability at least $1 - c_1 \exp(-c_2(d+1)\log p)$. *If* ρ_λ *is also* ℓ_0*-compatible, then with the same probability,*

$$\|\widetilde{B}(\widehat{\pi})\|_0 \lesssim r_{\min}(\Sigma) \frac{\lambda^2}{\rho'_\lambda(0+)^2} \|\widetilde{B}(\pi_0)\|_0. \tag{3.4}$$

The constants in (3.3) depend only on a_1 and a_2 in Conditions 3.1 and 3.2, while the constant in (3.4) also depends on $\overline{\rho}_0$ (from Definition 2.4). Both Theorems 3.1 and 3.2 are special cases of several general results which will be discussed further in Section 5.

Remark 3.1. Theorem 3.2 applies to regularizers that are not ℓ_0-compatible, such as the ℓ_1 penalty. For example, we can achieve $\|\widetilde{B}(\widehat{\pi})\|_1 \lesssim \|\widehat{B}\|_1 \lesssim \|\widetilde{B}(\pi_0)\|_1$ with the ℓ_1 penalty. The idea is that the penalty ρ_λ itself may be interpreted as a "measure of sparsity" that is weaker than the ℓ_0 norm.

Remark 3.2. Combining Theorems 3.1 and 3.2, with high probability,

$$\|\widehat{B} - \widetilde{B}(\widehat{\pi})\|_2^2 \lesssim \frac{\lambda^2}{r_{\min}(\Sigma)} s_0, \tag{3.5}$$

where $s_0 = \|\widetilde{B}(\pi_0)\|_0$ is the number of edges in a minimum-trace DAG. Let us investigate under which scalings of (n, p, d, s_0) the ℓ_2 error vanishes asymptotically. Choose λ to be on the order of $\sqrt{d \log p / n}$ and assume that Σ has bounded eigenvalues. Since the ℓ_2 error is essentially the sum over p regression problems, we normalize it by $1/p$. Thus, to achieve consistency in the normalized ℓ_2 error, it is sufficient to have $(s_0/p) d \log p \ll n$, where (s_0/p) is the average parent size of the minimum-trace DAG $\widetilde{B}(\pi_0)$. If $d \log p \leq \epsilon n$ with $\epsilon = o(1)$, then we can have $s_0/p \to \infty$ as long as $s_0/p \ll 1/\epsilon$, which imposes a quite weak sparsity assumption on $\widetilde{B}(\pi_0)$. Under this asymptotic scaling, Condition 3.1 allows $\tau_* \to 0$ for the MCP and ℓ_1, and Condition 3.2 allows $\rho_\lambda(\widetilde{B}(\pi_0))/\operatorname{tr} \widetilde{\Omega}(\pi_0) \to 0$. Furthermore, it is possible to allow $p \gg n$, which justifies our results for high-dimensional data. To establish normalized ℓ_2-consistency, van de Geer and Bühlmann (2013) assume that ϵ is sufficiently small but does not necessarily vanish asymptotically (their Condition 3.4). The fact that we need a slightly stronger assumption on d to obtain ℓ_2-consistency is probably the price we pay for obtaining uniform error control over all $\pi \in \mathbb{S}_p$.

Remark 3.3. Instead of normalizing the Frobenius error, we can also use (3.2) along with a similar argument as in Remark 3.2 in order to provide uniform control over the column deviations $\widehat{\beta}_j - \widetilde{\beta}_j(\widehat{\pi})$. It follows that $\sup_j \|\widehat{\beta}_j - \widetilde{\beta}_j(\widehat{\pi})\|_2 \to 0$ as long as $d \ll 1/\epsilon$, which highlights another advantage of obtaining uniform control over the neighbourhood problems.

3.3. **Discussion.** Let us discuss the assumptions required for Theorem 3.2 and compare the results derived so far to the existing literature.

3.3.1. *Conditions on* $\tau_*(\Sigma)$. Condition 3.1 on the signal strength, also called a *beta-min* condition in the literature, is only necessary in order to establish (3.4) in Theorem 3.2. Such a condition is not needed for the deviation bounds alone, or to establish the upper bound $\rho_\lambda(\widehat{B}) \lesssim \rho_\lambda(\widetilde{B}(\pi_0))$. Moreover, if we assume that the eigenvalues of Σ are bounded, then for penalties satisfying Condition 2.3 and for which $\rho'_\lambda(0+) = O(\lambda)$ (e.g., both the MCP and ℓ_1), Condition 3.1 holds if $\tau_* \gtrsim \lambda \gtrsim \sqrt{(d+1)\log p/n}$. This is the familiar scaling from the literature, up to a factor of $\sqrt{d}$ which reflects the cost of requiring *uniform* control over all possible neighbourhood regression problems. We have stated Condition 3.1 in terms of $\rho_\lambda(\tau_*)$ in order to allow for more general regularizers with arbitrary dependence on λ.

In the next section, we will extend these results to include model selection consistency. These results will be stated in terms of what we call *model selection exponents*, which are well-behaved as long as some kind of beta-min condition holds. In Section 4.4, we provide an example to illustrate that the usual beta-min condition required for model selection consistency is on the same order as the one required by Theorem 3.2, and thus we are able to guarantee consistent support recovery at no additional cost.

3.3.2. *Minimum-trace DAGs.* The use of a minimum-trace DAG in Theorem 3.2 leads one to ask several questions regarding such DAGs and how they relate to rest of $\mathfrak{D}(\Sigma)$. This is a fascinating and complex topic in and of itself, so we will only make a few comments on this matter here.

One may interpret the least squares loss in the following manner: Let $\ell(B, \Omega)$ denote the negative log-likelihood of (B, Ω). Then one may check that the following identity holds:

$$n^{-1}\ell(B, I) = \frac{1}{2n}\|\mathbf{X} - \mathbf{X}B\|_2^2 + \text{const.} \tag{3.6}$$

That is, when $\Omega = I$, the negative log-likelihood reduces (up to a constant) to the least squares loss. This implies that whenever there exists a DAG $\widetilde{B} \in \mathfrak{D}(\Sigma)$ with a corresponding variance matrix that is the identity, the estimator defined by (1.5) is in fact equivalent to a penalized maximum likelihood estimator for $\widetilde{B}$.

This fact has been exploited by van de Geer and Bühlmann (2013), who consider a similar (weighted) least squares estimator under the assumption that there exists an *equivariance* DAG, i.e. $\widetilde{B} \in \mathbb{D}_p$ such that (1.2) holds with

$\widetilde{\Omega} = \omega_0^2 I$. An equivariance DAG is known to be identifiable under certain assumptions (Peters and Bühlmann, 2013), whereas a minimum-trace DAG need not be identifiable. Moreover, it turns out that an equivariance DAG is automatically a minimum-trace DAG:

Lemma 3.1. *Suppose Σ is given and (1.2) holds for some $\widetilde{B} \in \mathbb{D}_p$ with $\widetilde{\Omega} = \omega_0^2 I$. Then $\widetilde{B} = \widetilde{B}(\pi_0)$ is the unique minimum-trace DAG.*

This is a simple consequence of the invariance of $\log\det\widetilde{\Omega}$ over the equivalence class $\mathfrak{D}(\Sigma)$ (see Appendix A.2). Thus, the minimum-trace property provides a natural generalization of the equivariance property to covariance matrices for which an equivariance DAG may not exist.

3.3.3. *An oracle inequality.* In order to control the sparsity of $\widehat{B}$ and $\widetilde{B}(\widehat{\pi})$, Theorem 3.2 requires a condition on a minimum-trace permutation π_0 (Condition 3.2). Here we provide a different kind of bound, in the form of an oracle inequality, which is true without any assumptions on π_0.

For any $\widetilde{B}(\pi) \in \mathfrak{D}(\Sigma)$, observe that $\frac{1}{2}\operatorname{tr}\widetilde{\Omega}(\pi)$ is the same as the expected loss function, i.e.

$$\mathbb{E}\Big[\frac{1}{2n}\|\mathbf{X} - \mathbf{X}\widetilde{B}(\pi)\|_2^2\Big] = \frac{1}{2}\operatorname{tr}\widetilde{\Omega}(\pi). \tag{3.7}$$

This gives us an expression for the expected penalized loss, which we denote

$$\widetilde{Q}_\lambda(\pi) := \frac{1}{2}\operatorname{tr}\widetilde{\Omega}(\pi) + \rho_\lambda(\widetilde{B}(\pi)).$$

The next result shows that the expected penalized loss of an estimated permutation $\widehat{\pi}$ is (up to a vanishing term) on the same order as that of the population-minimizer:

Lemma 3.2. *Suppose Conditions 2.1, 2.2, 2.3, and 3.1 hold. Then, there exist positive constants $\kappa_1 = \kappa_1(\Sigma)$ and $\kappa_2 = \kappa_2(\Sigma)$ such that if*

$$n > \kappa_1\, d\log p, \quad \lambda \geq \kappa_2\sqrt{\frac{(d+1)\log p}{n}},$$

then

$$\widetilde{Q}_\lambda(\widehat{\pi}) \;\leq\; \frac{3a_1+2}{a_1-2}\left(1 + 6\sqrt{\frac{(d+1)\log p}{n}}\right)\cdot \inf_\pi \widetilde{Q}_\lambda(\pi)$$

with probability at least $1 - c_1\exp(-c_2(d+1)\log p)$.

The proof of this inequality can be found in Appendix C.4. Note that the constant factor of $(3a_1+2)/(a_1-2)$ approaches 3 as the signal strength τ_* increases, and $\sqrt{(d+1)\log p/n} = o(1)$ as long as $n \gg (d+1)\log p$. If we let $\pi^* = \arg\min_\pi \widetilde{Q}(\pi)$, it follows from Lagrange duality that the DAG $\widetilde{B}(\pi^*) \in \mathfrak{D}(\Sigma)$ minimizes the weak sparsity measure ρ_λ subject to some upper bound on the expected prediction error, i.e., $\operatorname{tr}\widetilde{\Omega}(\pi) \leq c(\lambda)$. In this sense, the above oracle inequality gives another control on the sparsity of $\widetilde{B}(\widehat{\pi})$.

3.3.4. *Comparison.* The results in this section can be seen to generalize existing results from van de Geer and Bühlmann (2013) in several ways:

(1) We provide an extension of the penalized least-squares estimator to situations where the permutation π_0 may not be identifiable. Note that in order to exploit equivariance, one must essentially assume that the noise variance is known up to a multiplicative constant. By relaxing this assumption, our results provide theoretical guarantees in the practical setting where we have no prior knowledge about the error variances $\widetilde{\Omega}$.
(2) We allow for a general class of tractable penalties ρ_λ beyond the ℓ_0 penalty, which is known to be intractable. Although still nonconvex, with a concave penalty the program (1.5) is defined over a continuous parameter space and fast algorithms have been developed via block coordinate descent (e.g. Aragam and Zhou (2015)).
(3) We consider optimization over the full space of parameters $\mathbb{D}_p$, instead of a thresholded parameter space as in van de Geer and Bühlmann (2013) (see their Condition 3.3). This makes our method much more practical since the definition of the thresholded parameter space involves an unknown constant which is difficult to select.
(4) We have shown how $\widetilde{B}(\widehat{\pi})$ mimics the behaviour of an oracle minimizer of the expected penalized loss.
(5) Finally, in the next section, we show how $\widehat{B}$ attains exact support recovery when $p \gg n$, which has not been shown previously to the best of our knowledge. For a more detailed discussion see Remark 4.2.

Unlike van de Geer and Bühlmann (2013), however, due to our use of the LS loss instead of the log-likelihood, our results do not necessarily guarantee that π_0 is one of the sparsest permutations. Furthermore, our analysis stops short of providing a lower bound on $\|\widetilde{B}(\widehat{\pi})\|_0$, which allows for slightly stronger bounds compared to ours.

Finally, at a high-level, we note that our proof technique is completely different and makes an explicit connection between graphical models and existing concepts from the regression literature, and in this way sheds new light on the interplay between the covariance Σ and its regression coefficients through the intuitive concept of a neighbourhood regression problem. These ideas will be introduced in the next section.

4. Support recovery

The results in Section 3 provide some guarantees on the quality of the estimator $\widehat{B}$, but still fall short of establishing exact support recovery, i.e. $\operatorname{supp}(\widehat{B}) = \operatorname{supp}(\widetilde{B}(\widehat{\pi}))$. In this section we extend the results of the previous section and establish a general result that controls the probability $\mathbb{P}(\operatorname{supp}(\widehat{B}) \neq \operatorname{supp}(\widetilde{B}(\widehat{\pi})))$ in terms of so-called *neighbourhood regression problems*. The necessary material on neighbourhood regression will be introduced in Sections 4.1 and 4.2, which will allow us to state the main result

in Section 4.3. As stated, our main theorem is quite abstract, so we provide an example application of this theorem in Section 4.4. While the present section focuses on applying these results to the problem of support recovery, as the proofs in Section 5 will show, the material on neighbourhood regression will also be needed to establish the results from Section 3.

4.1. **Neighbourhood regression.** We begin with an intuitive description of what we mean by a "neighbourhood regression problem". In the literature on undirected graphical models, "the" neighbourhood of a fixed node X_j is implicitly understood to be the set of all other variables, namely the collection X_{-j}. For directed models we will need to generalize this notion of a neighbourhood to an arbitrary subset $S \subset [p]_j$ (which in particular includes the set $[p]_j$ corresponding to X_{-j}). At a high-level, the *neighbourhood regression problem* for the node X_j and a subset $S \subset [p]_j$ is the problem of estimating the linear projection of X_j onto X_S. These problems arise naturally from the iterative projection procedure described in Section 1.3: Doing this for all possible permutations is tantamount to projecting X_j onto all possible subsets $S \subset [p]_j$.

We now provide some formal definitions.

Definition 4.1. For any $S \subset [p]_j$, let

$$\beta_j(S) := \operatorname*{arg\,min}_{\beta \in \mathbb{R}^p,\ \mathrm{supp}(\beta) \subset S} \mathbb{E}\big[X_j - \beta^T X\big]^2.$$

We call $\beta_j(S)$ the SEM coefficients for variable j regressed on the variables S. The corresponding error (or noise) is given by the residual

$$\varepsilon_j(S) := X_j - \beta_j(S)^T X.$$

The support set of $\beta_j(S)$ is denoted by $m_j(S) := \mathrm{supp}(\beta_j(S))$ and the error variance by $\omega_j^2(S) := \mathrm{var}(\varepsilon_j(S))$.

All of the quantities in Definition 4.1 are *population* level quantities that depend on Σ, but not on the sample $\mathbf{X}$. We can now make explicit the connection between $\beta_j(S)$, $\widetilde{\beta}_j(\pi)$, and related quantities (recall the definition of $S_j(\pi)$ in (2.1)):

$$\begin{aligned} \widetilde{\beta}_j(\pi) &= \beta_j(S_j(\pi)), & \widetilde{\varepsilon}_j(\pi) &= \varepsilon_j(S_j(\pi)), \\ \mathrm{supp}(\widetilde{\beta}_j(\pi)) &= m_j(S_j(\pi)), & \widetilde{\omega}_j^2(\pi) &= \omega_j^2(S_j(\pi)). \end{aligned} \tag{4.1}$$

It is straightforward to verify each of these identities from the definition of $\widetilde{\beta}_j(\pi)$ as the jth column of $\widetilde{B}(\pi)$.

Now that we have defined the relevant population-level quantities, we turn to the problem of estimating these quantities via penalized least squares. We start with some fairly abstract notions in order to emphasize the key underlying assumptions.

Definition 4.2. Suppose $y \in \mathbb{R}^n$ and $Z \in \mathbb{R}^{n \times m}$. Let $S \subset [m]$ and consider the set defined by

$$\widehat{\Theta}_\lambda(y, Z; S) := \underset{\theta \in \mathbb{R}^m,\, \mathrm{supp}(\theta) \subset S}{\arg\min} \frac{1}{2n}\|y - Z\theta\|_2^2 + \rho_\lambda(\theta) \tag{4.2}$$

i.e., the set of global minimizers of the regularized, support-restricted, least-squares problem above. Let $\widehat{\Theta}_\lambda(y, Z) := \widehat{\Theta}_\lambda(y, Z; [m])$ correspond to the case where there is no support restriction.

Note that for the purposes of this abstract definition, y is considered a fixed quantity and may bear no relation to the matrix Z. Of course, in practice we are interested the case where y and Z are linked via a linear model:

Example 4.1. Consider the linear regression problem $\mathbf{y} = Z\theta^* + \mathbf{w}$, where $\mathbf{y} \in \mathbb{R}^n$, $Z \in \mathbb{R}^{n \times m}$, $\theta^* \in \mathbb{R}^m$ and $\mathbf{w} \sim \mathcal{N}_n(0, \sigma^2 I_n)$. Then $\widehat{\Theta}_\lambda(\mathbf{y}, Z)$ is the collection of penalized least squares estimators for θ^* in the classical linear regression set-up.

The support-restricted version $\widehat{\Theta}_\lambda(\mathbf{y}, Z; S)$ allows us to properly define a neighbourhood regression problem for the node X_j: A subset $S \subset [p]_j$ defines a neighbourhood of "candidate" regressors for X_j, and the neighbourhood regression problem for this node is given by $\widehat{\Theta}_\lambda(\mathbf{x}_j, \mathbf{X}; S)$, where $\mathbf{x}_j$ is the jth column of $\mathbf{X}$. Intuitively, it is the usual regression problem given the data $(y, Z) = (\mathbf{x}_j, \mathbf{X}_S)$.

Definition 4.3 (Neighbourhood regression). The *neighbourhood regression problem* for node X_j given a neighbourhood $S \subset [p]_j$ is defined to be the nonconvex program solved by $\widehat{\Theta}_\lambda(\mathbf{x}_j, \mathbf{X}; S)$. An arbitrary solution to this program will be denoted by $\widehat{\beta}_j(S)$, i.e. $\widehat{\beta}_j(S) \in \widehat{\Theta}_\lambda(\mathbf{x}_j, \mathbf{X}; S)$.

For any π such that $S_j(\pi) = S$, a solution $\widehat{\beta}_j(S)$ estimates $\widetilde{\beta}_j(\pi) = \beta_j(S)$ (by (4.1)). Thus neighborhood regression problems are the basic units in learning the parent set of a node and hence the DAG structure.

The concept of neighbourhood regression allows us to do two things:

(1) A key step in our proof is to show that the analysis of $\widehat{B}$ can be reduced to the study of a (very large) collection of neighbourhood regression problems;
(2) We will associate to each neighbourhood regression problem an "exponent", which is a fundamental quantity measuring how difficult the corresponding model selection problem is for a given regularizer. These exponents will then be used to write down explicit, non-asymptotic upper bounds on the model selection failure of $\widehat{B}$ in terms of the model selection failure of each neighbourhood regression (Proposition 5.5).

The details of these reductions can be found in Section 5.

4.2. **Model selection exponents.** Given some $n \times m$ matrix Z and m-vector θ^*, define a set of "bad" noise vectors as follows:

$$A(Z, \theta^*; S) := \Big\{ w \in \mathbb{R}^n : \mathrm{supp}(\widehat{\theta}) \neq \mathrm{supp}(\theta^*) \quad \exists\, \widehat{\theta} \in \widehat{\Theta}_\lambda(Z\theta^* + w, Z; S) \Big\}. \tag{4.3}$$

For a random vector $\mathbf{w} \in \mathbb{R}^n$ (e.g. $\mathbf{w} \sim \mathcal{N}_n(0, \sigma^2 I_n)$), we then have the following model selection failure event:

$$\mathcal{A}(\mathbf{w}, Z, \theta^*; S) := \Big\{ \mathbf{w} \in A(Z, \theta^*; S) \Big\}. \tag{4.4}$$

As usual we use the shorthand $\mathcal{A}(\mathbf{w}, Z, \theta^*) = \mathcal{A}(\mathbf{w}, Z, \theta^*; [m])$.

Definition 4.4. Given a regularizer ρ_λ, the *model selection exponent* for the regression problem $\mathbf{y} = Z\theta^* + \mathbf{w}$ is defined to be

$$\Phi_\lambda(Z, \theta^*, \sigma^2) := -\log \mathbb{P}\left[\mathcal{A}(\mathbf{w}, Z, \theta^*)\right],$$

where $\mathbb{P}$ is taken with respect to the distribution of $\mathbf{w} \sim \mathcal{N}_n(0, \sigma^2 I_n)$.

Remark 4.1. It is helpful to visualize the chain of dependencies that has led to Definition 4.4:

$$\rho_\lambda \to \widehat{\Theta}_\lambda(\mathbf{y}; Z) \to \mathcal{A}(\mathbf{w}, Z, \theta^*) \to \Phi_\lambda(Z, \theta^*, \sigma^2).$$

All of these quantities depend on ρ_λ, and hence also λ, even though this may be suppressed in the notation.

A larger exponent corresponds to better model selection performance. Let

$$\Psi_\lambda(\mathbf{X}, \Sigma) := \inf_{0 < \sigma \leq \sigma_{\max}} \inf_{\substack{\|\theta\|_0 \leq d(\Sigma) \\ \tau_*(\theta) \geq \tau_*(\Sigma)}} \Phi_\lambda(\mathbf{X},\, \theta,\, \sigma^2). \tag{4.5}$$

Recalling Definition 2.3, the function $\Psi_\lambda(\mathbf{X}, \Sigma)$ encodes what is usually proved in the regression literature: An upper bound on the probability of model selection failure given the maximum sparsity level $d(\Sigma)$, minimum signal strength $\tau_*(\Sigma)$, and the maximum variance $\sigma^2_{\max}$. This probability generally depends on λ, which in turn may depend on any of these quantities. In Section 4.4 we provide an example of such a bound under the MCP.

4.3. **Main result on model selection.** We are now in a position to strengthen our results by adding guarantees for model selection consistency.

Theorem 4.1. *Assume Condition 2.1. Then*

$$\mathbb{P}\left(\mathrm{supp}(\widehat{B}(\pi)) \neq \mathrm{supp}(\widetilde{B}(\pi)),\ \exists\, \pi \in \mathbb{S}_p \right) \ \leq\ p \binom{p}{d} \mathbb{E} e^{-\Psi_\lambda(\mathbf{X}, \Sigma)}.$$

Theorem 4.1 provides uniform control on the probability of false selection for all permutations $\pi \in \mathbb{S}_p$. We obtain model selection consistency for the

global minimizer $\widehat{B}$ as a special case: Apply this result to $\pi = \widehat{\pi}$ and note that $\widehat{B}(\widehat{\pi}) = \widehat{B}$. In addition, since $p\binom{p}{d} \leq p^{d+1}$, it follows that

$$\Psi_\lambda(\mathbf{X}, \Sigma) \geq C(d+1)\log p, \quad \text{for some } C > 1 \tag{4.6}$$

is sufficient for model selection consistency. Using existing results from the regression literature, a bound of the form (4.6) holds (i.e. with high probability for suitable regularizers when $\mathbf{X} \overset{\text{iid}}{\sim} \mathcal{N}_p(0, \Sigma)$), as illustrated in the next subsection.

Remark 4.2. van de Geer and Bühlmann (2013) remark that in a similar setting, a thresholded ℓ_0-penalized maximum likelihood estimator for (B, Ω) is model selection consistent if $s_0 \log p/n \to 0$, where s_0 is the number of edges in a so-called *minimal I-map* (i.e. a DAG $\widetilde{B}(\pi)$ with the fewest edges). This result allows for p to grow with n, but is not truly high-dimensional: In the general case, $s_0 > p$ in which case this requires $p \ll n$. By contrast, our results allow $\|\widetilde{B}(\pi_0)\|_0$ as large as $O(pd)$ as long as $d \log p/n \to 0$, in which case Theorem 4.1 guarantees model selection consistency when $p \gg n$.

4.4. **Example.** Recall the MCP (Example 2.1) and suppose $\mathbf{y} = Z\theta^* + \mathbf{w}$. Huang et al. (2012) consider a penalized least squares estimator as in (4.2), and provide conditions under which global minima of the resulting program are model selection consistent. In particular, we have the following result which bounds the model selection exponents under the MCP:

Lemma 4.1. *Take $\rho_\lambda = \rho_\lambda(\cdot\,; \gamma)$ as in (2.7) and assume Σ is positive definite with bounded eigenvalues. Assume that*

(1) $d(\Sigma) \leq \kappa_4 \cdot \min\{p, n, n/\log p\}$,

(2) $\tau_*(\Sigma) > (1+\gamma)\lambda$ for some $\gamma > \kappa_5 > 0$.

Then for any $\lambda \geq \kappa_6 \cdot \sqrt{(d+1)\log p/n}$, it follows that $\mathbb{E}e^{-\Psi_\lambda(\mathbf{X},\Sigma)} \leq 3\exp\left(-2\min\{d\log p, n\}\right)$. Here, $\kappa_j = \kappa_j(\Sigma)$ $(j = 4, 5, 6)$ are constants depending only on $\{r_{\min}(\Sigma), r_{\max}(\Sigma)\}$.

This result, whose proof is omitted, is a consequence of Theorem 4.2 in Huang et al. (2012) and Proposition 2 in Zhang and Huang (2008), along with Proposition B.3, which is needed for the case $\beta_j(S) = 0$.

Taking λ as in Theorem 3.1 and combining Lemma 4.1 with Theorem 4.1, it follows that as long as $n \geq C(d+1)\log p$ for a sufficiently large constant $C > 0$,

$$\mathbb{P}\left(\operatorname{supp}(\widehat{B}) \neq \operatorname{supp}(\widetilde{B}(\widehat{\pi})),\ \exists\, \widehat{\pi} \in \widehat{\mathbb{S}}_p\right) \to 0,$$

which proves the model selection consistency of $\widehat{B}$ under the MCP. We reiterate that the sparsity of $\widetilde{B}(\widehat{\pi})$ is controlled by that of $\widetilde{B}(\pi_0)$ via Theorem 3.2.

5. Proofs

The proof of Theorems 3.1, 3.2, and 4.1 will be broken down into three main steps. First, we establish some basic properties of the objective function and the probability space in order to provide a uniform bound on the probability of false selection for any neighbourhood problem (Proposition 5.5). Then, we will derive an independent result that gives a deviation bound for a *fixed* design regression problem (Theorem 5.1). Using similar arguments from the first step, we then show that this result can be applied uniformly to each neighbourhood problem, which yields our claimed deviation bounds (see Proposition B.1). Finally, the last step is to use a "basic inequality" argument to control the sparsity.

5.1. Support recovery.

We begin by controlling the event

$$\mathcal{B} := \{\text{supp}(\widehat{B}(\pi)) \neq \text{supp}(\widetilde{B}(\pi)) \ \exists\, \pi \in \mathbb{S}_p\}. \tag{5.1}$$

We will do this by reducing the analysis of $\widehat{B}(\pi)$ to a series of neighbourhood regression problems. There are two key steps: (i) Showing that each estimator $\widehat{B}(\pi)$ is equivalent to solving a series of p regression problems given by $\widehat{\Theta}_\lambda(\mathbf{x}_j, \mathbf{X}; S)$, and (ii) Controlling the total number of sets S that need to be considered.

5.1.1. *Reduction to neighbourhood regression.* Recall that the jth column of $\widehat{B}(\pi)$ is denoted by $\widehat{\beta}_j(\pi)$ and as usual, denote the sample version of $\widetilde{\varepsilon}_j(\pi)$ by boldface, i.e. $\widetilde{\boldsymbol{\varepsilon}}_j(\pi) := \mathbf{x}_j - \mathbf{X}\widetilde{\beta}_j(\pi)$. The first step above is justified by the following result. The symbol $\perp\!\!\!\perp$ is used to denote independence of random variables.

Lemma 5.1. *Suppose* $\mathbf{X} \overset{iid}{\sim} \mathcal{N}_p(0, \Sigma)$ *and* $\lambda \geq 0$. *Then the following statements are true:*

(a) *For any* $j \in [p]$ *and* $\pi \in \mathbb{S}_p$, $\widetilde{\boldsymbol{\varepsilon}}_j(\pi) \perp\!\!\!\perp \mathbf{X}_{S_j(\pi)}$.

(b) *A matrix* $\widehat{B}(\pi) \in \mathbb{D}_p$ *is a restricted global minimizer (Definition 2.1) if and only if* $\widehat{\beta}_j(\pi) \in \widehat{\Theta}_\lambda(\mathbf{x}_j, \mathbf{X}; S_j(\pi))$ *for each* $j \in [p]$.

(c) $\widehat{B} = \widehat{B}(\widehat{\pi})$ *is a global minimizer of* $Q(B)$ *if and only if* $\widehat{\beta}_j(\widehat{\pi}) \in \widehat{\Theta}_\lambda(\mathbf{x}_j, \mathbf{X}; S_j(\widehat{\pi}))$ *for each* $j \in [p]$ *and* $\widehat{\pi} \in \widehat{\mathbb{S}}_p$.

The proof of this lemma, which is a simple consequence of how the least squares loss and the regularizer factor, is found in Appendix A.3. This allows us to formally establish the equivalence between the DAG problem and neighbourhood regression: In order to construct $\widehat{B}(\pi)$, it suffices to solve a neighbourhood regression problem for each column of $\widehat{B}(\pi)$, given by $\widehat{\Theta}_\lambda(\mathbf{x}_j, \mathbf{X}; S_j(\pi))$. A key observation is that through the independence established in Lemma 5.1(a) and a conditioning argument, we can reduce the regression problem given by $\widehat{\Theta}_\lambda(\mathbf{x}_j, \mathbf{X}; S_j(\pi))$ to a fixed design problem. The details are outlined in the proof of Proposition 5.5.

5.1.2. *Invariant sets and monotonicity.* As a consequence of Lemma 5.1, we have (cf. (5.1))

$$\mathcal{B} = \bigcup_{j=1}^{p} \{\text{supp}(\widehat{\beta}_j(\pi)) \neq \text{supp}(\widetilde{\beta}_j(\pi)) \ \exists\, \pi \in \mathbb{S}_p\}.$$

Since there are $p!$ total permutations, in principle the event $\{\text{supp}(\widehat{\beta}_j(\pi)) \neq \text{supp}(\widetilde{\beta}_j(\pi)) \exists\, \pi \in \mathbb{S}_p\}$ involves controlling a superexponential number of estimators, which seems hopeless. In order to reduce the total number of estimators we must control, we introduce the notion of an *invariant set*:

Definition 5.1. For any $S \subset [p]_j$, define a collection of subsets by

$$\mathcal{T}_j(S) := \{T \subset [p]_j : \ \beta_j(T) = \beta_j(S)\} = \{T \subset [p]_j : \ m_j(T) = m_j(S)\},$$

where $\beta_j(S)$ and $m_j(S)$ are defined in Definition 4.1. If $T \in \mathcal{T}_j(S)$, we call T an *invariant set of S for j*, or *S-invariant* for short.

In other words, for any j, $\mathcal{T}_j(S)$ is the collection of candidate sets $T \subset [p]_j$ such that the projection of X_j onto $\{X_i, i \in T\}$ is invariant. With some abuse of terminology, let us refer to $m_j(T) = \text{supp}(\beta_j(T))$ as the *support of neighbourhood* T (for node j). An equivalent description of $\mathcal{T}_j(S)$ is the set of neighbourhoods T whose support (for node j) is the same and equals $m_j(S)$.

Example 5.1. Continuing with Example 1.1, one can verify that $\mathcal{T}_3(\{1,2\}) = \{\{1\}, \{1,2\}\}$ and $m_3(\{1,2\}) = \{1\}$. Note also that $\mathcal{T}_3(S) = \mathcal{T}_3(\{1,2\})$ for any $S \in \mathcal{T}_3(\{1,2\})$. In the same example,

$$\{\{1\}, \{1,2\}\} = \mathcal{T}_3(\{1,2\}) \neq \mathcal{T}_3(\{1,4\}) = \{\{1,4\}, \{1,2,4\}\}.$$

The following lemma illustrates a crucial property of invariant sets:

Lemma 5.2. $T_1, T_2 \in \mathcal{T}_j(S) \implies T_1 \cup T_2 \in \mathcal{T}_j(S)$.

In other words, if two neighbourhoods share the same support, the union of these neighbourhoods must also have the same support. This justifies the following definition:

Definition 5.2. The unique largest element of $\mathcal{T}_j(S)$ shall be denoted by $M_j(S)$. Formally,

$$M_j(S) := \bigcup \mathcal{T}_j(S) = \bigcup \{T \subset [p]_j : \ \beta_j(T) = \beta_j(S)\}.$$

For instance, in Example 5.1 we have $M_3(\{1\}) = \{1,2\}$. The name "S-invariant set" comes from the fact that for any $T \in \mathcal{T}_j(S)$, we have the following useful identities:

$$\beta_j(m_j(S)) = \beta_j(S) = \beta_j(T) = \beta_j(M_j(S)), \tag{5.2}$$

$$\varepsilon_j(m_j(S)) = \varepsilon_j(S) = \varepsilon_j(T) = \varepsilon_j(M_j(S)). \tag{5.3}$$

The reason for introducing invariant sets is that it is generally *sufficient* to study the neighbourhood problem for $M_j(S)$ in the sense that once we

have model selection consistency for each estimator in $\widehat{\Theta}_\lambda(\mathbf{x}_j, \mathbf{X}; M_j(S))$, the same is *guaranteed* for estimators based on every other neighbourhood in $\mathcal{T}_j(S)$. In fact, we have the following result, which says that model selection properties of the S-restricted estimators are monotone with respect to those sets S that contain the true support.

Lemma 5.3. *Suppose that $Z \in \mathbb{R}^{n\times m}$ is fixed and consider the regression problem $\mathbf{y} = Z\theta^* + \mathbf{w}$ for some $\theta^* \in \mathbb{R}^m$. If $\mathrm{supp}(\theta^*) \subset S \subset U$, then we have the following inclusion: $A(Z, \theta^*; S) \subset A(Z, \theta^*; U)$. In particular,*

$$\mathcal{A}(\mathbf{w}, Z, \theta^*; S) \subset \mathcal{A}(\mathbf{w}, Z, \theta^*; U), \tag{5.4}$$

where $A(Z, \theta^*; S)$ and $\mathcal{A}(\mathbf{w}, Z, \theta^*; S)$ are defined in (4.3)–(4.4).

Intuitively, for a fixed support, the set of "bad" noise vectors for the larger problem involving U is at least as big as the set of "bad" noise vectors for the smaller problem involving S. We are interested in the model selection failure of $\widehat{\beta}_j(S)$ for $\beta_j(S)$, which can be stated as

$$\Big\{ \mathrm{supp}(\widehat{\beta}_j(S)) \neq \mathrm{supp}(\beta_j(S)) \\ \text{for some } \widehat{\beta}_j(S) \in \widehat{\Theta}_\lambda(\mathbf{x}_j, \mathbf{X}; S) \Big\} = \mathcal{A}(\widetilde{\boldsymbol{\varepsilon}}_j(S), \mathbf{X}, \beta_j(S); S) \tag{5.5}$$

in the notation introduced in (4.4). The next result encapsulates a notion of *monotonicity* that is used throughout the rest of the proof:

Corollary 5.4. *Suppose $\mathbf{X} \overset{iid}{\sim} \mathcal{N}_p(0, \Sigma)$. For any $S \subset [p]_j$, we have*

$$\mathcal{A}\Big(\widetilde{\boldsymbol{\varepsilon}}_j(S), \mathbf{X}, \beta_j(S); S\Big) \subset \mathcal{A}\Big(\widetilde{\boldsymbol{\varepsilon}}_j(M_j(S)), \mathbf{X}, \beta_j(M_j(S)); M_j(S)\Big). \tag{5.6}$$

In other words, the invariant sets for any fixed neighbourhood are monotone in the same sense as Lemma 5.3. Thus, in order to control the neighbourhood problem for some set S, it suffices to control the strictly harder problem given by $M_j(S)$.

5.1.3. *A bound on false selection.* For any $\Sigma \succ 0$ and fixed node X_j, define the following collections of subsets:

$$m_j(\Sigma) := \{m_j(S) : \ S \subset [p]_j\}, \tag{5.7}$$

$$M_j(\Sigma) := \{M_j(S) : \ S \subset [p]_j\}. \tag{5.8}$$

Note that $|m_j(\Sigma)| = |M_j(\Sigma)|$. As long as it is clear whether the argument is a set S or a matrix Σ, this should not cause any confusion with $m_j(S)$ and $M_j(S)$. In Example 1.1, $m_3(\Sigma) = \{\emptyset, \{1\}, \{2\}, \{4\}, \{1, 4\}, \{2, 4\}\}$. In contrast, one can verify that $m_2(\Sigma) = \{\emptyset, \{1\}, \{3\}, \{4\}, \{3, 4\}\}$ which illustrates how some nodes have a reduced set of parent sets.

For any neighbourhood $S \subset [p]_j$, recall that the associated error variance is given by $\omega_j^2(S) = \mathrm{var}(\varepsilon_j(S))$ (Definition 4.1). With some more abuse of notation, let

$$\Phi_j(S) := \Phi_\lambda(\mathbf{X}_S, \, (\beta_j(S))_S, \, \omega_j^2(S)). \tag{5.9}$$

Note that we must restrict the SEM coefficients $\beta_j(S)$ to the subset S in order for this exponent to be well-defined. Since $\mathrm{supp}(\beta_j(S)) \subset S$, this does not change anything.

We have the following general result, whose proof can be found in Appendix A.7:

Proposition 5.5. *Fix $j \in [p]$. Under Condition 2.1, we have*

$$\mathbb{P}\left(\mathrm{supp}(\widehat{\beta}_j(S)) \neq \mathrm{supp}(\beta_j(S)),\ \exists\, S \subset [p]_j\right) \ \leq \sum_{T \in m_j(\Sigma)} \mathbb{E} e^{-\Phi_j(M_j(T))},$$

where $m_j(\Sigma)$ is defined by (5.7) and $\Phi_j(\,\cdot\,)$ is defined by (5.9).

Proposition 5.5 says that to control the probability of false selection uniformly for all neighbourhoods S of the node j, it suffices to control a much smaller class of problems given by the neighbhourhoods $M_j(T)$ for each support set $T \in m_j(\Sigma)$. By Definition 2.3, $|m_j(S)| \leq d$ for all j and S, which implies that $|m_j(\Sigma)| \leq \binom{p}{d} \leq p^d$. Since there are 2^{p-1} subsets of $[p]_j$, whenever d is much smaller than p the cardinality of $m_j(\Sigma)$ is much smaller than that of $2^{[p]_j}$.

5.1.4. *Proof of Theorem 4.1.* For any $T \in m_j(\Sigma)$, Lemma 5.3 applied with $S = M_j(T)$ and $U = [p]$ yields

$$\Phi_j(M_j(T)) \geq \Phi_\lambda(\mathbf{X},\, \beta_j(T),\, \omega_j^2(T)).$$

Recalling $d(\Sigma)$ and $\tau_*(\Sigma)$ in Definition 2.3, we have $\|\beta_j(T)\|_0 \leq d(\Sigma)$ and $\tau_*(\beta_j(T)) \geq \tau_*(\Sigma)$, as well as $\omega_j^2(T) \leq \sigma_{\max}^2$. The previous expression combined with (4.5) implies:

$$\Phi_j(M_j(T)) \geq \Psi_\lambda(\mathbf{X}, \Sigma) \quad \text{for all } T \in M_j(\Sigma). \tag{5.10}$$

Combining Proposition 5.5, (5.10) and a union bound over $j \in [p]$,

$$\begin{aligned}
&\mathbb{P}\left(\mathrm{supp}(\widehat{\beta}_j(S)) \neq \mathrm{supp}(\beta_j(S)),\ \exists\, j \in [p],\, S \subset [p]_j\right) \\
&\qquad \leq \sum_{j=1}^{p} \sum_{T \in m_j(\Sigma)} \mathbb{E} \exp(-\Phi_j(M_j(T)) \\
&\qquad \leq p\binom{p}{d}\, \mathbb{E} \exp(-\Psi_\lambda(\mathbf{X}, \Sigma)), \qquad (5.11)
\end{aligned}$$

since there are at most $\binom{p}{d}$ subsets in $m_j(\Sigma)$. In order to complete the proof, note that Lemma 5.1 implies

$$\begin{aligned}
&\left\{ \mathrm{supp}(\widehat{B}(\pi)) \neq \mathrm{supp}(\widetilde{B}(\pi)),\ \exists\, \pi \in \mathbb{S}_p \right\} \\
&\qquad = \left\{ \mathrm{supp}(\widehat{\beta}_j(\pi)) \neq \mathrm{supp}(\widetilde{\beta}_j(\pi)),\ \exists\, \pi \in \mathbb{S}_p,\, j \in [p] \right\} \\
&\qquad = \left\{ \mathrm{supp}(\widehat{\beta}_j(S)) \neq \mathrm{supp}(\beta_j(S)),\ \exists\, S \subset [p]_j,\, j \in [p] \right\}.
\end{aligned}$$

In the last line we used $\widetilde{\beta}_j(\pi) = \beta_j(S_j(\pi))$. Combined with (5.11), this gives the desired result.

5.2. **Deviation bounds.** In order to establish uniform control over the probability of false selection for all possible neighbourhood regression problems in the previous section, we relied on a monotonicity property (cf. Lemma 5.3) of model selection. Unfortunately, this may not hold for weaker types of consistency. In order to bound the deviations $\|\widehat{B}(\pi) - \widetilde{B}(\pi)\|_r$ $(r = 1, 2)$, we will use a modified argument that invokes a different kind of monotone class.

5.2.1. *An upper bound for fixed design.* We start by establishing a general bound on the ℓ_r $(r = 1, 2)$ estimation errors for a fixed design regression problem with a general regularizer ρ_λ. The objective here is to derive conditions under which we can guarantee such bounds for a fixed design problem, and then show that these conditions hold uniformly for all neighbourhood problems. The conditions we will need are familiar from the literature: A *Gaussian width condition* and a *restricted eigenvalue condition*.

For the rest of this subsection, we let $Z \in \mathbb{R}^{n\times m}$ and $w \in \mathbb{R}^n$ be a fixed matrix and fixed vector, respectively.

Definition 5.3 (Gaussian width). We say that the *Gaussian width (GW) condition* holds for (w, Z) relative to ρ_λ if there is a numerical constant $\delta \in (0, 1)$ such that

$$\frac{1}{n}|\langle w, Zu\rangle| \leq \delta\left[\frac{1}{2n}\|Zu\|_2^2 + \rho_\lambda(u)\right], \ \forall u \in \mathbb{R}^m,$$

in which case we write $(w, Z) \in \mathrm{GW}_{\rho_\lambda}(\delta)$. If this inequality is strict for all $u \neq 0$, we write $(w, Z) \in \mathrm{GW}^\circ_{\rho_\lambda}(\delta)$.

We will be interested in the case where both w and Z are allowed to be random but independent. In this setting, for Gaussian designs considered in this paper, the GW condition holds with high probability for the ℓ_1 penalty (this follows from a standard Hölder inequality argument), and has similarly been shown to hold for penalties induced by ℓ_q norms for $0 \leq q \leq 1$ (Raskutti et al., 2011). Zhang and Zhang (2012) provide a version of this condition that applies to general nonconvex regularizers.

Before we proceed, let us note the following key relation between model selection consistency and the GW condition:

Lemma 5.6. *Consider the setup of Lemma 5.3, namely, the regression problem* $\mathbf{y} = Z\theta^* + \mathbf{w}$ *but with* $\theta^* = 0$. *Then*

$$\mathcal{A}(\mathbf{w}/\delta,\, Z,\, 0)^c = \big\{(\mathbf{w}, Z) \in \mathrm{GW}^\circ_{\rho_\lambda}(\delta)\big\}. \tag{5.12}$$

Thus, in order to ensure the GW condition for $(\mathbf{w}, Z_S)$, it suffices to show that the corresponding regression problem is model selection consistent when the *true coefficients are all set to zero* and the noise variance is inflated by

a factor of $1/\delta^2$. Zhang and Zhang (2012) refer to this property as *null-consistency*.

For any set $A \subset [m]$ and $\xi > 0$, define the following "cone":

$$C_{\rho_\lambda}(A, \xi) := \{u \in \mathbb{R}^m : \rho_\lambda(u_{A^c}) \leq \xi \rho_\lambda(u_A)\}. \tag{5.13}$$

This definition also depends on the ambient dimension m; when we wish to emphasize this we will write $C^m_{\rho_\lambda}(A, \xi)$. The term "cone" here is used in an extended sense, in analogy with the ℓ_1 cone found in previous work.

Definition 5.4 (Generalized restricted eigenvalue). The *generalized restricted eigenvalue (RE) constant* of Z with respect to ρ_λ over a subset A is

$$\phi^2_{\rho_\lambda}(Z, A;\, \xi) := \inf \left\{ \frac{\|Zu\|_2^2}{n\|u\|_2^2} : u \in C_{\rho_\lambda}(A, \xi),\, u \neq 0 \right\}. \tag{5.14}$$

In the sequel, we often suppress the dependence of the generalized RE constants on λ and ξ, writing $\phi^2_\rho(Z, A) = \phi^2_{\rho_\lambda}(Z, A; \xi)$. Note that the usual restricted eigenvalue is equivalent to the special case $\rho_\lambda = \lambda\|\cdot\|_1$ (Bickel et al., 2009).

Consider the usual linear regression set up, $y = Z\theta^* + w$, where $\theta^* \in \mathbb{R}^m$ and we define $S^* = \mathrm{supp}(\theta^*)$. The following general result establishes that the two conditions $(w, Z) \in \mathrm{GW}_\rho(\delta)$ and $\phi^2_\rho(Z, S^*) > 0$ are sufficient to bound the deviation $\widehat{\theta} - \theta^*$:

Theorem 5.1. *Assume $(w, Z) \in \mathrm{GW}_{\rho_\lambda}(\delta)$ for some ρ_λ satisfying Condition 2.2 and $\delta \in (0, 1)$. Let $\xi = \xi(\delta) := (1+\delta)/(1-\delta)$ and assume $\phi^2 := \phi^2_\rho(Z, S^*; \xi) > 0$. Then any $\widehat{\theta} \in \widehat{\Theta}_\lambda(Z\theta^* + w, Z)$ satisfies*

$$\|\widehat{\theta} - \theta^*\|_2 \leq \frac{2\,\xi}{\phi^2}\, \rho'_\lambda(0+) \|\theta^*\|_0^{1/2}, \tag{5.15}$$

$$\|\widehat{\theta} - \theta^*\|_1 \leq \frac{2\,\xi(1+\xi)}{\phi^2}\, \rho'_\lambda(0+) \|\theta^*\|_0. \tag{5.16}$$

The proof of Theorem 5.1 is found in Appendix A.8. The GW condition is quantified by the constant $\delta \in (0, 1)$, and the restricted eigenvalue condition depends on the free parameter $\xi > 0$; these two are linked via the relation $\xi(\delta) = (1+\delta)/(1-\delta)$ and play subtle roles in the proof.

5.2.2. *Uniform deviation bounds.* In order to apply Theorem 5.1 to the neighbourhood problems for Σ, we first provide a uniform bound on the restricted eigenvalues $\phi^2_\rho(\mathbf{X}_S, m_j(S))$ in terms of the smallest eigenvalue of Σ. More precisely, we show in Proposition B.4 that with high probability $\phi^2_\rho(\mathbf{X}_S, A) \geq r_{\min}(\Sigma)$ for all pairs $A \subset S$ with $|A| \leq d$.

The next step is to show that with high probability, the GW condition holds for $(\widetilde{\boldsymbol{\varepsilon}}_j(S), \mathbf{X}_S)$ for all j and S. First, let us define

$$\mathcal{E}_S(\delta, \lambda) = \mathcal{E}_S(\delta, \lambda; j) := \left\{ \left(\widetilde{\boldsymbol{\varepsilon}}_j(S), \mathbf{X}_S\right) \in \mathrm{GW}^\circ_{\rho_\lambda}(\delta) \right\} \tag{5.17}$$

According to Lemma 5.6, we have $\mathcal{E}_S(\delta,\lambda) = \mathcal{A}\big(\widetilde{\varepsilon}_j(S)/\delta, \mathbf{X}, 0; S\big)^c$. An immediate consequence is that the complements of the events $\mathcal{E}_S(\delta,\lambda)$ are monotonic in the sense that they obey Corollary 5.4 when $\beta_j(S)$ is replaced with the zero vector. Analogous to the arguments leading up to Proposition 5.5 and its proof, this allows us to provide uniform control on $\mathcal{E}_S(\delta,\lambda)$, in a sense made precise by Proposition B.3. Once again, the key is to reduce the total number of events to control—*a priori* superexponential in size—down to a tractable number which can be controlled.

Together, Propositions B.3 and B.4 show that we have uniform control over both the restricted eigenvalues and the Gaussian widths for the neighbourhood problems $\widehat{\Theta}_\lambda(\mathbf{x}_j, \mathbf{X}; S)$. Thus, we can apply Theorem 5.1 to each of these problems and obtain, with high probability, deviation bounds of the form

$$\|\widehat{\beta}_j(S) - \beta_j(S)\|_2 \le \frac{2\,\xi}{r_{\min}(\Sigma)} \rho'_\lambda(0+) \cdot \|\beta_j(S)\|_0^{1/2}$$

for all j and $S \subset [p]_j$. The precise statement can be found in Proposition B.1. Theorem 3.1 is an immediate consequence of these two results, the details of which are discussed in Appendix B, along with the statement and proof of the results mentioned in this subsection.

5.3. **Sparsity.** The last step is to control the sparsity of the estimate $\widehat{B}$ and the candidate DAG $\widetilde{B}(\widehat{\pi})$. We start with a kind of basic inequality that is adapted to the present, nonidentifiable setting:

Lemma 5.7. *Let* $\mathbf{E}(\pi) := \mathbf{X} - \mathbf{X}\widetilde{B}(\pi)$. *For any* $\pi \in \mathbb{S}_p$ *and* $\widehat{\pi} \in \widehat{\mathbb{S}}_p$,

$$\begin{aligned} \frac{1}{2n}\|\mathbf{X}(\widetilde{B}(\widehat{\pi}) - \widehat{B})\|_2^2 + \rho_\lambda(\widehat{B}) \;\le\; & \frac{1}{2n}\|\mathbf{E}(\pi)\|_2^2 - \frac{1}{2n}\|\mathbf{E}(\widehat{\pi})\|_2^2 \\ & + \frac{1}{n}\operatorname{tr}\Big(\mathbf{E}(\widehat{\pi})^T\mathbf{X}(\widetilde{B}(\widehat{\pi}) - \widehat{B})\Big) \\ & + \rho_\lambda(\widetilde{B}(\pi)). \end{aligned} \tag{5.18}$$

In contrast to the basic inequality used in usual regression, (5.18) also leverages the many possible decompositions of $\mathbf{X}$, one for each permutation. Three terms in this inequality are particularly important:

(1) The difference in residuals $\|\mathbf{E}(\pi)\|_2^2/(2n) - \|\mathbf{E}(\widehat{\pi})\|_2^2/(2n)$ explains the origin of the minimum-trace permutation: We would like to make $\|\mathbf{E}(\pi)\|_2^2/(2n)$ as small as possible in order to minimize this difference. By standard concentration arguments, $\|\mathbf{E}(\pi)\|_2^2/n$ is close to its expectation, $\operatorname{tr}\widetilde{\Omega}(\pi)$ (cf. (3.7)). Hence, we choose π to minimize $\operatorname{tr}\widetilde{\Omega}(\pi)$.

(2) The quantity $\operatorname{tr}(\mathbf{E}(\widehat{\pi})^T\mathbf{X}(\widetilde{B}(\widehat{\pi}) - \widehat{B}))/n$ can be bounded using the Gaussian width condition (Definition 5.3). There is a subtlety regarding whether to decompose this along rows or columns; see Lemma C.1.

(3) The penalty on $\widehat{B}$ can be replaced with $\rho_\lambda(\widetilde{B}(\widehat{\pi}))$ by showing that $\rho_\lambda(\widehat{B}) \gtrsim \rho_\lambda(\widetilde{B}(\widehat{\pi}))$. This is true whenever the deviation $\widehat{B} - \widetilde{B}(\widehat{\pi})$ is not too large, which is guaranteed by Theorem 3.1.

Once we have established control of these three terms, it is not hard to show that $\rho_\lambda(\widetilde{B}(\widehat{\pi})) \lesssim \rho_\lambda(\widehat{B}) \lesssim \rho_\lambda(\widetilde{B}(\pi_0))$. More precisely, we have the following bound in terms of the constants δ (cf. Definition 5.3) and a_2 (cf. Condition 3.2):

$$\frac{2\delta}{1-\delta}\rho_\lambda(\widetilde{B}(\widehat{\pi})) \leq \rho_\lambda(\widehat{B}) \leq \frac{2}{1-\delta}\Big(1 + \frac{10}{a_2}\Big)\rho_\lambda(\widetilde{B}(\pi_0)).$$

The precise statement and proof are given by Proposition C.3. Theorem 3.2 follows as a consequence of this bound and an ℓ_0-compatibility argument, the details of which can be found in Appendix C.

6. Discussions

To conclude, we discuss various issues related to extending these results derived herein to more general settings. We begin by addressing some technical details involved with the ℓ_0 regularizer (see also Remark 2.1), which is followed by a discussion of some extensions of our results to more general graphical models.

6.1. **The ℓ_0 case.** Condition 2.2 precludes discontinuous regularizers such as the ℓ_0 penalty, however, a simple modification to the proofs accounts for any regularizer that is ℓ_0-compatible as defined in Definition 2.4, including the ℓ_0 penalty (Example 2.3). All of the proofs go through as-is, except the proof of Theorem 5.1, which requires a minor modification (see Remark A.1 in the Appendix). The modified ℓ_2 bound becomes

$$\|\widehat{\theta} - \theta^*\|_2 \leq \sqrt{\frac{2\overline{\rho}_0\xi}{\phi^2}} \cdot \lambda\|\theta^*\|_0^{1/2}.$$

The proof of Lemma C.2 also simplifies significantly. Thus everything derived here applies equally to the ℓ_0 penalty. In particular, Theorem 3.2 implies $\|\widetilde{B}(\widehat{\pi})\|_0 \lesssim \|\widehat{B}\|_0 \lesssim \|\widetilde{B}(\pi_0)\|_0$. If $\widetilde{B}(\pi_0)$ is one of the sparsest DAGs in $\mathfrak{D}(\Sigma)$, then the number of edges in $\widehat{B}$ will be on the same order as that of the sparsest ones. For example, this occurs when an equivariance DAG is also amongst the sparsest, as assumed by van de Geer and Bühlmann (2013).

6.2. **Undirected models and extensions.** As an immediate corollary of the results derived in this work, we also obtain consistency of the popular neighbourhood-regression based estimator of the Gaussian graphical model introduced in Meinshausen and Bühlmann (2006). This follows since, in particular, we have uniform control over the neighbourhood problems

$\widehat{\Theta}_\lambda(\mathbf{x}_j, \mathbf{X}; [p]_j)$ for each $j = 1, \dots, p$. In this case, there are only p neighborhood problems to control, much smaller than $p\binom{p}{d}$ for DAGs. One advantage of our results is the relaxation of the strong *neighbourhood stability* assumption (also known as the *irrepresentability condition*) through the use of nonconvex penalties such as the MCP and SCAD.

Furthermore, our results can also be applied to cyclic graphs (also known as *non-recursive SEM*). This is a consequence of the fact that our proof technique essentially circumvents the acyclicity constraint by providing uniform control over all possible neighbourhood problems. The decision to present our results in the concrete setting of acyclic models was entirely in order to present the main ideas in a familiar setting and to avoid unnecessary technical digressions and generalities.

Appendix A. Proofs of technical results

A.1. **Proof of Lemma 1.1.** We need the following simple lemma, which follows since $P_\pi A = PAP^T$ for some permutation matrix P:

Lemma A.1. $A = MNM^T \iff P_\pi A = (P_\pi M)(P_\pi N)(P_\pi M)^T$.

Recall the modified Cholesky decomposition of A (also called the LDLT decomposition): $A = LDL^T$ for a lower triangular matrix L, with unit diagonal entries, and a diagonal matrix D. When A is positive definite, the pair (L, D) is unique and we refer to it as the *Cholesky decomposition of* A.

Note that $(\widetilde{B}, \widetilde{\Omega}) \in \mathbb{D}_p \times \mathbb{R}^p_+$ satisfies (1.2) if and only if $\Sigma^{-1} = (I - \widetilde{B})\widetilde{\Omega}^{-1}(I - \widetilde{B}^T)$. This is because a zero-mean Gaussian distribution is uniquely determined by its covariance. Let us denote the set of all such pairs $(\widetilde{B}, \widetilde{\Omega})$ as $\mathfrak{D}'$. Next, note that $\widetilde{B} \in \mathbb{D}_p$ if and only if $P_\pi \widetilde{B}$ is lower triangular for some permutation π. Lemma A.1 implies that $(\widetilde{B}, \widetilde{\Omega}) \in \mathfrak{D}'$ iff $(I - P_\pi \widetilde{B}, P_\pi \widetilde{\Omega}^{-1})$ is a Cholesky decomposition of $P_\pi \Sigma^{-1}$ for some π.

By the definition (1.9), $(I - P_\pi \widetilde{B}(\pi), P_\pi \widetilde{\Omega}(\pi)^{-1})$ is also a Cholesky decomposition of $P_\pi \Sigma^{-1}$. Since the Cholesky decomposition is unique for positive definite matrices, we have $(\widetilde{B}, \widetilde{\Omega}) \in \mathfrak{D}'$ iff $(\widetilde{B}, \widetilde{\Omega}) = (\widetilde{B}(\pi), \widetilde{\Omega}(\pi))$ for some π, which gives the desired result, since $\mathfrak{D}(\Sigma)$ is the projection of $\mathfrak{D}'$ onto its first coordinate.

A.2. **Proof of Lemma 3.1.** Consider the following program:

$$\min \sum_{j=1}^p x_j^2 \text{ subject to } \sum_{j=1}^p \log x_j^2 = C. \tag{A.1}$$

The solution to this program is given by $x_j^2 = e^{C/p}$ for all $j = 1, \dots, p$. In other words, the minimum is attained by a constant vector. It is straightforward to verify that $\log\det\widetilde{\Omega}(\pi) = \log\det\Sigma$ and hence $\log\det\widetilde{\Omega}(\pi) = \sum_j \log\widetilde{\omega}_j^2(\pi)$ is constant for all $\pi \in \mathbb{S}_p$. Thus for any $\pi \in \mathbb{S}_p$, the vector $(\widetilde{\omega}_1^2(\pi), \dots, \widetilde{\omega}_p^2(\pi)) \in \mathbb{R}^p$ is feasible for (A.1), which implies that $\operatorname{tr}\widetilde{\Omega}(\pi)$ is

minimized whenever $\widetilde{\omega}_1^2(\pi) = \cdots = \widetilde{\omega}_p^2(\pi)$. Finally, uniqueness of $\widetilde{B}(\pi_0)$ follows from the uniqueness of the Cholesky decomposition.

A.3. **Proof of Lemma 5.1.** The first conclusion (a) follows from elementary properties of conditional expectation and the identity

$$\mathbb{E}(X_j \,|\, X_{S_j(\pi)}) = \widetilde{\beta}_j(\pi)^T X.$$

Since (c) is a special case of (b), it suffices to prove (b).

Fix $\pi \in \mathbb{S}_p$ and let $S_j = S_j(\pi)$. If $\widehat{\beta}_j(\pi) \in \widehat{\Theta}_\lambda(\mathbf{x}_j, \mathbf{X}; S_j)$ for each j, then evidently $\widehat{B}(\pi) = [\,\widehat{\beta}_1(\pi) \,|\, \cdots \,|\, \widehat{\beta}_p(\pi)\,]$ minimizes $Q(B)$ over $\mathbb{D}_p[\pi]$. For the reverse direction, recall that $\mathbf{X}_{S_j}$ is the $n \times |S_j|$ matrix formed by extracting the columns in S_j, and similarly for $(\beta_j)_{S_j}$. For any $B \in \mathbb{D}_p[\pi]$ we have $(\beta_j)_{S_j^c} = 0$ for each j, so we can write

$$\begin{aligned}\frac{1}{2n}\|\mathbf{X} - \mathbf{X}B\|_2^2 + \rho_\lambda(B) &= \sum_{j=1}^p \Big\{ \frac{1}{2n}\|\mathbf{x}_j - \mathbf{X}\beta_j\|_2^2 + \rho_\lambda(\beta_j) \Big\} \\ &= \sum_{j=1}^p \Big\{ \frac{1}{2n}\|\mathbf{x}_j - \mathbf{X}_{S_j}(\beta_j)_{S_j}\|_2^2 + \rho_\lambda((\beta_j)_{S_j}) \Big\}.\end{aligned}$$

Then $\widehat{B}(\pi) = [\,\widehat{\beta}_1(\pi) \,|\, \cdots \,|\, \widehat{\beta}_p(\pi)\,] \in \min_{\mathbb{D}_p[\pi]} Q(B)$ if and only if

$$\widehat{\beta}_j(\pi) \in \arg\min_\beta \frac{1}{2n}\|\mathbf{x}_j - \mathbf{X}\beta\|_2^2 + \rho_\lambda(\beta) \quad \text{subject to } \beta_{S_j^c} = 0.$$

In other words, $\widehat{\beta}_j(\pi) \in \widehat{\Theta}_\lambda(\mathbf{x}_j, \mathbf{X}; S_j)$ for each j. Since π was arbitrary, the desired claim follows.

A.4. **Proof of Lemma 5.2.** The proof relies on the following property of L^2 projections: For any two sets $S, R \subset [p]_j$, we have

$$\beta_j(S \cup R) = \beta_j(S) \iff \varepsilon_j(S) \perp X_i,\ \forall i \in R. \tag{A.2}$$

To lighten the notation, let $S^* = m_j(S)$. Note that $\beta_j(S) = \beta_j(S^*)$ since $\operatorname{supp}(\beta_j(S)) = S^*$. It follows from (A.2) that $\varepsilon_j(S^*) \perp\!\!\!\perp X_i$ for $i \in S \setminus S^*$. Similarly, since $\operatorname{supp}(\beta_j(T_k)) = S^*$, we have $\varepsilon_j(S^*) \perp\!\!\!\perp X_i$ for $i \in T_k \setminus S^*$ and $k = 1, 2$. It follows that

$$\varepsilon_j(S^*) \perp\!\!\!\perp X_i,\ \forall i \in (T_1 \setminus S^*) \cup (T_2 \setminus S^*)$$

hence the application of (A.2) in the reverse direction yields

$$\beta_j(T_1 \cup T_2) = \beta_j\big(S^* \cup (T_1 \setminus S^*) \cup (T_2 \setminus S^*)\big) = \beta_j(S^*) = \beta_j(S).$$

A.5. **Proof of Lemma 5.3.** It suffices to show

$$A(Z,\theta^*;U)^c \subset A(Z,\theta^*;S)^c.$$

Suppose $w \in A(Z,\theta^*;U)^c$, i.e., $\operatorname{supp}(\widetilde{\theta}) = \operatorname{supp}(\theta^*) := S^*$ for any $\widetilde{\theta} \in \widehat{\Theta}_\lambda(Z\theta^* + w, Z; U)$. We wish to show that for any $\widehat{\theta} \in \widehat{\Theta}_\lambda(Z\theta^* + w, Z; S)$, it must also be true that $\operatorname{supp}(\widehat{\theta}) = S^*$. Let

$$F(\theta) = \frac{1}{2n}\|Z(\theta^* - \theta) + w\|_2^2 + \rho_\lambda(\theta)$$

denote the objective function in Definition 4.2 of $\widehat{\Theta}_\lambda(y, Z; S)$ with $y = Z\theta^* + w$. Since $\operatorname{supp}(\widehat{\theta}) \subset S \subset U$, $\widehat{\theta}$ is feasible for the U-restricted problem, whence

$$F(\widetilde{\theta}) \leq F(\widehat{\theta})$$

for any $\widetilde{\theta} \in \widehat{\Theta}_\lambda(Z\theta^* + w, Z; U)$. But $\widetilde{\theta}$ is also feasible for the S-restricted problem since $\operatorname{supp}(\widetilde{\theta}) = S^* \subset S$, so that

$$F(\widetilde{\theta}) \geq F(\widehat{\theta}) \implies F(\widetilde{\theta}) = F(\widehat{\theta}).$$

Since the value $F(\widetilde{\theta})$ is by definition the global minimum of F for the U-restricted problem and $\operatorname{supp}(\widehat{\theta}) \subset U$, $\widehat{\theta}$ must be a global minimizer of F for the U-restricted problem, i.e., $\widehat{\theta} \in \widehat{\Theta}_\lambda(Z\theta^* + w, Z; U)$, whence $\operatorname{supp}(\widehat{\theta}) = S^*$ as desired.

A.6. **Proof of Corollary 5.4.** By Lemma 5.3 and the fact that $S \subset M_j(S)$, we have

$$\mathcal{A}\Big(\widetilde{\boldsymbol{\varepsilon}}_j(S), \mathbf{X}, \beta_j(S); S\Big) \subset \mathcal{A}\Big(\widetilde{\boldsymbol{\varepsilon}}_j(S), \mathbf{X}, \beta_j(S); M_j(S)\Big). \tag{A.3}$$

Using (5.2) and (5.3), we have the following identity:

$$\mathcal{A}\Big(\widetilde{\boldsymbol{\varepsilon}}_j(S), \mathbf{X}, \beta_j(S); M_j(S)\Big) = \mathcal{A}\Big(\widetilde{\boldsymbol{\varepsilon}}_j(M_j(S)), \mathbf{X}, \beta_j(M_j(S)); M_j(S)\Big).$$

Plugging this into (A.3) yields the desired result.

A.7. **Proof of Proposition 5.5.** Throughout, for simplicity, let

$$\mathcal{A}_S := \mathcal{A}(\widetilde{\boldsymbol{\varepsilon}}_j(S), \mathbf{X}, \beta_j(S); S).$$

Fix $S \subset [p]_j$ and let $\theta^* = \beta_j(S)$, $s^* = |m_j(S)| = \|\theta^*\|_0$ and $\varepsilon^* = \widetilde{\boldsymbol{\varepsilon}}_j(S)$ so that $\mathcal{A}_S = \mathcal{A}(\varepsilon^*, \mathbf{X}, \theta^*; S)$. Note that $\mathcal{A}(\varepsilon^*, \mathbf{X}, \theta^*; S)$ represents the following model selection failure:

$$\operatorname{supp}(\widehat{\theta}) \neq \operatorname{supp}(\theta^*) \quad \exists\, \widehat{\theta} \in \widehat{\Theta}_\lambda(\mathbf{X}\theta^* + \varepsilon^*, \mathbf{X}; S).$$

Since $\operatorname{supp}(\theta^*) \subset S$, we can restrict $\mathbf{X}$ and θ^* to S, so that the above is equivalent to

$$\operatorname{supp}(\widehat{\theta}) \neq \operatorname{supp}(\theta^*_S) \quad \exists\, \widehat{\theta} \in \widehat{\Theta}_\lambda(\mathbf{X}_S\theta^*_S + \varepsilon^*, \mathbf{X}_S).$$

which is the same event as $\mathcal{A}(\varepsilon^*, \mathbf{X}_S, \theta^*_S)$. To summarize, $\mathcal{A}_S = \mathcal{A}(\varepsilon^*, \mathbf{X}_S, \theta^*_S)$.

Since ε^* is independent of $\mathbf{X}_S$ by Lemma 5.1(a), by conditioning on $\mathbf{X}_S$ we are dealing with a fixed design regression problem with Gaussian noise $\varepsilon^* = \widetilde{\boldsymbol{\varepsilon}}_j(S) \sim \mathcal{N}_n(0, \omega_j^2(S) I_n)$. We obtain

$$\begin{aligned}\mathbb{P}(\mathcal{A}_S) &= \mathbb{E}\Big[\,\mathbb{P}\left(\mathcal{A}(\varepsilon^*, \mathbf{X}_S, \theta^*_S)\right) \mid \mathbf{X}_S\Big] \\ &\le \mathbb{E}\exp[-\Phi_\lambda(\mathbf{X}_S, \theta^*_S, \omega_j^2(S))] \\ &= \mathbb{E}\exp(-\Phi_j(S)),\end{aligned}$$

where the last line uses (5.9). Now we have

$$\{\mathrm{supp}(\widehat{\beta}_j(S)) \neq \mathrm{supp}(\beta_j(S)),\ \exists\, S \subset [p]_j\} = \bigcup_{S \subset [p]_j} \mathcal{A}_S \ \subset \bigcup_{T \in m_j(\Sigma)} \mathcal{A}_{M_j(T)},$$

where the equality is by (5.5) and the inclusion follows from Corollary 5.4. Note that this is the key step where the reduction occurs. Hence, we have

$$\begin{aligned}\mathbb{P}\Big(\bigcup_{S \subset [p]_j} \mathcal{A}_S\Big) &\le \mathbb{P}\Big(\bigcup_{T \in m_j(\Sigma)} \mathcal{A}_{M_j(T)}\Big) \\ &\le \sum_{T \in m_j(\Sigma)} \mathbb{P}(\mathcal{A}_{M_j(T)}) \le \sum_{T \in m_j(\Sigma)} \mathbb{E}\exp(-\Phi_j(M_j(T))),\end{aligned}$$

which is the desired probability bound.

A.8. **Proof of Theorem 5.1.** Recall that $S^* := \mathrm{supp}(\theta^*)$. To lighten notation, for any vector u let $u_1 := u_{S^*}$, $u_2 := u_{(S^*)^c}$, and also $\Delta := \widehat{\theta} - \theta^*$. Then invoking the subadditivity of ρ_λ (this is a consequence of Condition 2.2),

$$\begin{aligned}\rho_\lambda(\widehat{\theta}) - \rho_\lambda(\theta^*) &= \rho_\lambda(\Delta + \theta^*) - \rho_\lambda(\theta^*) \\ &= \rho_\lambda(\Delta_1 + \theta^*_1) + \rho_\lambda(\Delta_2) - \rho_\lambda(\theta^*_1) \\ &\ge -\rho_\lambda(\Delta_1) + \rho_\lambda(\Delta_2). \qquad \text{(A.4)}\end{aligned}$$

It is straightforward to derive

$$\frac{1}{2n}\|y - Z\widehat{\theta}\|_2^2 - \frac{1}{2n}\|y - Z\theta^*\|_2^2 = \frac{1}{2n}\|Z\Delta\|^2 - \frac{1}{n}\langle w, Z\Delta\rangle. \qquad \text{(A.5)}$$

Since $(w, Z) \in \mathrm{GW}_{\rho_\lambda}(\delta)$, we can invoke the GW condition with $u = \Delta$,

$$-\frac{1}{n}\langle w, Z\Delta\rangle \ge -\frac{1}{n}|\langle w, Z\Delta\rangle| \ge -\delta\frac{1}{2n}\|Z\Delta\|^2 - \delta\rho_\lambda(\Delta). \qquad \text{(A.6)}$$

It follows that

$$\begin{aligned}
0 &\geq \frac{1}{2n}\|y - Z\widehat{\theta}\|_2^2 - \frac{1}{2n}\|y - Z\theta^*\|_2^2 + \rho_\lambda(\widehat{\theta}) - \rho_\lambda(\theta^*) \\
&\geq \frac{1}{2n}\|Z\Delta\|^2 - \frac{1}{n}\langle w, Z\Delta\rangle - \rho_\lambda(\Delta_1) + \rho_\lambda(\Delta_2) \\
&\geq \frac{1-\delta}{2n}\|Z\Delta\|^2 - \delta\rho_\lambda(\Delta) - \rho_\lambda(\Delta_1) + \rho_\lambda(\Delta_2) \\
&= \frac{1-\delta}{2n}\|Z\Delta\|^2 - (1+\delta)\rho_\lambda(\Delta_1) + (1-\delta)\rho_\lambda(\Delta_2) \\
&= (1-\delta)\Big[\frac{1}{2n}\|Z\Delta\|^2 + \rho_\lambda(\Delta_2) - \xi\rho_\lambda(\Delta_1)\Big], \qquad \text{(A.7)}
\end{aligned}$$

where the first inequality by optimality of $\widehat{\theta}$, the second by (A.5), and the third by (A.6). The next line follows from an an application of $\rho_\lambda(\Delta) = \rho_\lambda(\Delta_1) + \rho_\lambda(\Delta_2)$. Since $\delta < 1$ by assumption, it follows that $\rho_\lambda(\Delta_2) \leq \xi\rho_\lambda(\Delta_1)$ which implies $\Delta \in C_\rho(S^*, \xi(\delta))$.

Recalling the definition (5.14) of $\phi_\rho^2(Z, S^*)$, we conclude that $\frac{1}{2n}\|Z\Delta\|^2 \geq \frac{\phi^2}{2}\|\Delta\|_2^2$ which combined with (A.7), dropping $\rho_\lambda(\Delta_2)$, gives

$$0 \;\geq\; \frac{\phi^2}{2}\|\Delta\|_2^2 - \xi\rho_\lambda(\Delta_1).$$

Combining with the following (note $\|\Delta_1\|_0 \leq \|\theta^*\|_0$),

$$\rho_\lambda(\Delta_1) \leq \rho'_\lambda(0+)\|\Delta_1\|_1 \leq \rho'_\lambda(0+)\|\theta^*\|_0^{1/2}\|\Delta\|_2 \qquad \text{(A.8)}$$

and re-arranging proves (5.15). For (5.16), since $\Delta \in C_\rho(S^*, \xi(\delta))$, we can use Lemma B.7 to construct a set $M \subset [p]$ with $|M| = |S^*| = \|\theta^*\|_0$ such that $\Delta \in C_1(M, \xi(\delta))$. Then

$$\begin{aligned}
\|\Delta\|_1 = \|\Delta_M\|_1 + \|\Delta_{M^c}\|_1 &\leq (1+\xi)\|\Delta_M\|_1 \\
&\leq (1+\xi)\|\theta^*\|_0^{1/2}\|\Delta_M\|_2 \\
&\leq \frac{2\,\xi(1+\xi)}{\phi^2}\cdot\rho'_\lambda(0+)\|\theta^*\|_0.
\end{aligned}$$

Remark A.1. Theorem 5.1 also applies to penalties that are ℓ_0-compatible, even if Condition 2.2 does not hold. Simply replace (A.8) with

$$\rho_\lambda(\Delta_1) \leq \overline{\rho}_0\lambda^2\|\Delta_1\|_0 \leq \overline{\rho}_0\lambda^2\|\theta^*\|_0.$$

The rest of the proof requires no changes.

A.9. **Proof of Lemma 5.6.** If $(\mathbf{w}, Z) \in \mathrm{GW}^\circ_{\rho_\lambda}(\delta)$, then for any $u \neq 0$,

$$\begin{aligned}
&\frac{\delta}{2n}\|Zu\|_2^2 - \frac{1}{n}\mathbf{w}^T Zu + \delta\rho_\lambda(u) > 0 \\
\iff\; &\frac{1}{2n}\|\mathbf{w}/\delta - Zu\|_2^2 + \rho_\lambda(u) > \frac{1}{2n}\|\mathbf{w}/\delta\|_2^2.
\end{aligned}$$

The latter inequality implies

$$\{0\} = \arg\min_u \|\mathbf{w}/\delta - Zu\|_2^2/(2n) + \rho_\lambda(u),$$

that is, 0 is the unique global minimizer of the right hand side. Recalling the definition of $\mathcal{A}\big(\mathbf{w}/\delta, Z, 0\big)$ in (4.4), we obtain the desired result.

A.10. **Proof of Lemma 5.7.** Observe that for any $\pi \in \mathbb{S}_p$,

$$Q(\widehat{B}) \le Q(\widehat{B}(\pi)) \le Q(\widetilde{B}(\pi)), \tag{A.9}$$

where $\widehat{B}(\pi)$ is a restricted global minimizer as in (2.2). Moreover, we have the following alternative expression for Q:

$$Q(B) = \frac{1}{2n}\|\mathbf{X}(\widetilde{B}(\widehat{\pi}) - B) + \mathbf{E}(\widehat{\pi})\|_2^2 + \rho_\lambda(B), \quad \text{for any } \widehat{\pi} \in \widehat{\mathbb{S}}_p. \tag{A.10}$$

Thus, using (A.9) and (A.10),

$$\begin{aligned}
0 &\le Q(\widetilde{B}(\pi)) - Q(\widehat{B}) \\
&= \frac{1}{2n}\|\mathbf{E}(\pi)\|_2^2 - \frac{1}{2n}\|\mathbf{X}(\widetilde{B}(\widehat{\pi}) - \widehat{B}) - \mathbf{E}(\widehat{\pi})\|_2^2 + \rho_\lambda(\widetilde{B}(\pi)) - \rho_\lambda(\widehat{B}) \\
&= \frac{1}{2n}\|\mathbf{E}(\pi)\|_2^2 - \frac{1}{2n}\|\mathbf{E}(\widehat{\pi})\|_2^2 - \frac{1}{2n}\|\mathbf{X}(\widetilde{B}(\widehat{\pi}) - \widehat{B})\|_2^2 \\
&\quad + \frac{1}{n}\operatorname{tr}\Big(\mathbf{E}(\widehat{\pi})^T\mathbf{X}(\widetilde{B}(\widehat{\pi}) - \widehat{B})\Big) + \rho_\lambda(\widetilde{B}(\pi)) - \rho_\lambda(\widehat{B}).
\end{aligned}$$

Since (A.9) holds for any π, this completes the proof.

Appendix B. Auxiliary results for deviation bounds

This section contains the bulk of our technical results on controlling the deviations $\widehat{\beta}_j(S) - \beta_j(S)$ for a neighbourhood problem. These constitute the main ingredients used in proving Theorem 3.1, which we also do in this section.

For any $\delta \in (0,1)$ and $\lambda \ge 0$, define the following events:

$$\mathcal{E}(\delta,\lambda) = \Big\{\big(\widetilde{\boldsymbol{\varepsilon}}_j(S), \mathbf{X}_S\big) \in \mathrm{GW}^\circ_{\rho_\lambda}(\delta),\ \forall j \in [p], S \subset [p]_j\Big\}, \tag{B.1}$$

$$\mathcal{R}(\delta) = \Big\{\phi^2_\rho(\mathbf{X}_S, m_j(S)) \ge r_{\min}(\Sigma) > 0,\ \forall j \in [p], S \subset [p]_j\Big\}. \tag{B.2}$$

Note that $\mathcal{E}(\delta,\lambda) = \bigcap_{j=1}^p \bigcap_{S\subset[p]_j} \mathcal{E}_S(\delta,\lambda;j)$, where $\mathcal{E}_S(\delta,\lambda;j)$ was defined in (5.17).

The structure of the proofs will be to show that on one or both of these events, the desired conclusions follow. Explicit bounds on the probabilities of these events are established in Section B.1.

B.1. **Uniform deviation bounds.** As discussed in Section 5.2.2, one of the main ingredients in proving Theorem 3.1 is a general result regarding deviation bounds for neighbourhood problems, given by Proposition B.1 below. The other ingredient is Proposition B.2 below, regarding the behavior of certain model selection exponents defined in analogy with (4.5):

$$\psi_\lambda(\mathbf{X}, \sigma^2_{\max}; \delta) := \inf_{0 \le \sigma \le \sigma_{\max}} \Phi_\lambda(\mathbf{X},\, 0,\, \sigma^2/\delta^2). \tag{B.3}$$

We often suppress the dependence on δ and write $\psi_\lambda(\mathbf{X}, \sigma^2_{\max})$. Note that, in view of Lemma 5.6, $\psi_\lambda(\mathbf{X}, \sigma^2_{\max})$ describes the conditional probability, given $\mathbf{X}$, that $(\sigma\mathbf{w}, \mathbf{X})$ violates a GW condition, where $\mathbf{w} \sim \mathcal{N}_n(0, I_n)$ is independent of $\mathbf{X}$. More precisely,

$$\begin{aligned}\sup_{0\le\sigma\le\sigma_{\max}} \mathbb{P}\left[(\sigma\mathbf{w}, \mathbf{X}) \notin \mathrm{GW}^\circ_{\rho_\lambda}(\delta) \;\middle|\; \mathbf{X}\right] &= \sup_{0\le\sigma\le\sigma_{\max}} \exp[-\Phi_\lambda(\mathbf{X},\, 0,\, \sigma^2/\delta^2)] \\ &= \exp[-\psi_\lambda(\mathbf{X}, \sigma^2_{\max})].\end{aligned}$$

We also recall the relation

$$\xi = \xi(\delta) = \frac{1+\delta}{1-\delta}. \tag{B.4}$$

Proposition B.1. *Assume that Σ satisfies Condition 2.1 and ρ_λ satisfies Condition 2.2. Suppose $\mathbf{X} \overset{iid}{\sim} \mathcal{N}_p(0, \Sigma)$, $\delta \in (0,1)$, and define ξ by (B.4). Then there exist constants $c_0, c_1, c_2 > 0$ such that the following holds: If*

$$n > c_0 \frac{\sigma^2_{\max}(1+\xi)^2}{r_{\min}(\Sigma)} d \log p,$$

then with probability at least $1 - c_1 \exp(-c_2 n) - p\binom{p}{d}\mathbb{E}\exp(-\psi_\lambda(\mathbf{X}, \sigma^2_{\max}; \delta))$,

$$\|\widehat{\beta}_j(S) - \beta_j(S)\|_2 \le \frac{2\,\xi}{r_{\min}(\Sigma)} \rho'_\lambda(0+) \cdot \|\beta_j(S)\|_0^{1/2}, \tag{B.5}$$

$$\|\widehat{\beta}_j(S) - \beta_j(S)\|_1 \le \frac{2\,\xi(1+\xi)}{r_{\min}(\Sigma)} \rho'_\lambda(0+) \cdot \|\beta_j(S)\|_0, \tag{B.6}$$

uniformly over all $j \in [p]$ and $S \subset [p]_j$.

For future reference, inspection of the proof shows that the conclusion of Proposition B.1 holds on $\mathcal{E}(\delta, \lambda) \cap \mathcal{R}(\delta)$. For regularizers that satisfy Condition 2.3 we have the following control on the exponent $\psi_\lambda(\mathbf{X}, \sigma^2_{\max})$:

Proposition B.2. *Assume that $\mathbf{X} \overset{iid}{\sim} \mathcal{N}_p(0, \Sigma)$, and that ρ_λ satisfies Condition 2.3. Then there exist constants $c > 0$ and $C = C(\underline{\rho}_1, \underline{\rho}_0)$ such that for any $\delta \in (0,1)$, if*

$$\lambda \ge C\delta^{-1}\sigma_{\max}\|\Sigma\|^{1/4}\sqrt{\frac{(d+1)\log p}{n}} \tag{B.7}$$

then $\mathbb{E}\exp(-\psi_\lambda(\mathbf{X}, \sigma^2_{\max}; \delta)) \le c\exp(-\min\{2(d+1)\log p, n\})$.

The proof of Proposition B.2 follows from an argument similar to that in Zhang and Zhang (2012) and is omitted for brevity. In order to prove Proposition B.1, we need the following two intermediate results, providing uniform control on RE constants and GW conditions. Recall $\mathcal{E}(\delta, \lambda)$ as defined in (B.1).

Proposition B.3 (Uniform GW control)**.** *For any $\delta \in (0,1)$ and $\lambda > 0$,*

$$\mathbb{P}[\mathcal{E}(\delta,\lambda)] \geq 1 - p\binom{p}{d}\mathbb{E}\exp\left[-\psi_\lambda(\mathbf{X}, \sigma^2_{\max}; \delta)\right].$$

Proof. Fix $\delta \in (0,1)$. By analogy with (5.9), for any neighbourhood $S \subset [p]_j$, let

$$\xi_j(S) \;:=\; \Phi_\lambda(\mathbf{X}_S,\, 0,\, \omega_j^2(S)/\delta^2) \;\geq\; \psi_\lambda(\mathbf{X}, \sigma^2_{\max}; \delta), \tag{B.8}$$

where the inequality follows from (B.3) and the definition of $\sigma_{\max}$ in (2.3). We follow the proof of Proposition 5.5, but with $\beta_j(S)$ replaced with 0, and $\widetilde{\varepsilon}_j(S)$ replaced with $\widetilde{\varepsilon}_j(S)/\delta$. To simplify, let $\mathcal{F}_S^j = \mathcal{E}_S(\delta,\lambda;j)^c$ and $\mathcal{E} = \mathcal{E}(\delta,\lambda)$. By the comment following (5.17),

$$\mathcal{F}_S^j = \mathcal{A}\big(\widetilde{\varepsilon}_j(S)/\delta, \mathbf{X}, 0; S\big) = \mathcal{A}\big(\widetilde{\varepsilon}_j(S)/\delta, \mathbf{X}_S, 0\big)$$

where the second equality is by the same argument in the proof of Proposition 5.5. Since $\widetilde{\varepsilon}_j(S)/\delta \sim \mathcal{N}\big(0, [\omega_j^2(S)/\delta^2] I_n\big)$ independent of $\mathbf{X}_S$, we conclude, using Definition 4.4, that

$$\mathbb{P}\left(\mathcal{F}_S^j \mid \mathbf{X}_S\right) = \exp[-\xi_j(S)],$$

hence $\mathbb{P}(\mathcal{F}_S^j) \leq \mathbb{E}\exp[-\psi_\lambda(\mathbf{X}, \sigma^2_{\max})]$, $\forall S \subset [p]_j$, using the inequality in (B.8). The events $\mathcal{F}_S^j$ are monotonic in S according to Corollary 5.4. (The division of $\varepsilon_j(S)$ by δ does not change anything in that proof.) It follows that

$$\mathcal{E}^c = \bigcup_{j=1}^{p} \bigcup_{S \subset [p]_j} \mathcal{F}_S^j \;\subset\; \bigcup_{j=1}^{p} \bigcup_{T \in m_j(\Sigma)} \mathcal{F}_{M_j(T)}^j.$$

Taking the union bound, and using $|m_j(\Sigma)| \leq \binom{p}{d}$ and

$$\mathbb{P}\left[\,\mathcal{F}_{M_j(T)}^j\,\right] \leq \mathbb{E}\exp[-\psi_\lambda(\mathbf{X}, \sigma^2_{\max})], \quad \forall\, T \in m_j(\Sigma),$$

finishes the proof. □

Proposition B.4 (Uniform RE control)**.** *Assume $\mathbf{X} \overset{iid}{\sim} \mathcal{N}_p(0,\Sigma)$, Σ satisfies Condition 2.1 and ρ_λ satisfies Condition 2.2. There exist universal constants $c_0, c_1, c_2 > 0$, such that if*

$$n > c_0\, \frac{\sigma^2_{\max}(1+\xi)^2}{r_{\min}(\Sigma)} d(\Sigma) \log p$$

then with probability at least $1 - c_1\exp(-c_2 n)$,

$$\inf_{1\leq j\leq p} \inf_{S\subset[p]_j} \inf_{\substack{A\subset S\\ |A|\leq d}} \phi^2_\rho(\mathbf{X}_S, A;\, \xi) \;\geq\; r_{\min}(\Sigma).$$

The proof of this proposition appears in Section B.3 below. Recalling the definition of $\mathcal{R}(\delta)$ in (B.2), combined with $m_j(S) = \|\beta_j(S)\|_0 \le d$ (cf. Definition 2.3), Proposition B.4 implies that $\mathcal{R}(\delta)$ holds with probability at least $1 - c_1 \exp(-c_2 n)$. Let us show how Proposition B.1 follows.

Proof of Proposition B.1. Recall the definitions of $\mathcal{E}(\delta, \lambda)$ in (B.1) and $\mathcal{R}(\delta)$ in (B.2). Propositions B.3 and B.4 guarantee that

$$\mathbb{P}\left(\mathcal{R}(\delta) \cap \mathcal{E}(\delta, \lambda)\right) \ge 1 - c_1 \exp(-c_2 n) - p\binom{p}{d}\mathbb{E}\exp(-\psi_\lambda(\mathbf{X}, \sigma^2_{\max}; \delta)).$$

Thus, it suffices to deduce (B.5) and (B.6) whenever we are on the event $\mathcal{R}(\delta) \cap \mathcal{E}(\delta, \lambda)$. The case $\beta_j(S) = 0$ follows from (5.17) and Lemma 5.6, and the case $\beta_j(S) \neq 0$ follows from Theorem 5.1 applied to the corresponding neighbourhood regression problems. □

B.2. **Proof of Theorem 3.1.** Let us consider the ℓ_2 bound, noting that the ℓ_1 version follows similarly:

$$\|\widehat{B} - \widetilde{B}(\widehat{\pi})\|_2 \le \frac{2\xi}{r_{\min}(\Sigma)} \rho'_\lambda(0+) \|\widetilde{B}(\widehat{\pi})\|_0^{1/2}. \tag{B.9}$$

(B.9) can be seen by applying Proposition B.1 with $S = S_j(\widehat{\pi})$, and noting that $\widehat{\beta}_j(\widehat{\pi}) = \widehat{\beta}_j(S_j(\widehat{\pi}))$, $\widetilde{\beta}_j(\widehat{\pi}) = \widetilde{\beta}_j(S_j(\widehat{\pi}))$, and the fact that $\|u_j - v_j\|_2 \le a\|v_j\|_0^{1/2}$ for all j, implies $\sum_j \|u_j - v_j\|_2^2 \le a^2 \sum_j \|v_j\|_0$. The bound on the probability of (B.9) follows by combining Proposition B.1 with Proposition B.2.

B.3. **Uniform control of restricted eigenvalues.** In this section, we collect the necessary results that lead to the proof of Proposition B.4. We begin with a definition which generalizes the familiar (ℓ_1) restricted eigenvalue (Bickel et al., 2009; Raskutti et al., 2011):

Definition B.1. $Z \in \mathbb{R}^{n \times m}$ satisfies a *generalized restricted eigenvalue condition of order k w.r.t. ρ_λ* with parameters $\alpha, \xi > 0$, denoted as $Z \in \mathrm{RE}_{\rho_\lambda}(k, \alpha; \xi)$, if

$$\frac{1}{n}\|Zu\|_2^2 \ge \alpha^2 \|u\|_2^2 \quad \forall u \in C_{\rho_\lambda}(A, \xi),$$

uniformly for all $A \subset [m]$ with $|A| = k$. Equivalently, recalling Definition 5.4,

$$Z \in \mathrm{RE}_{\rho_\lambda}(k, \alpha; \xi) \iff \inf_{\substack{A \subset [m] \\ |A| = k}} \phi^2_{\rho_\lambda}(Z, A, \xi) \ge \alpha^2.$$

In the sequel, we will suppress the dependence of various quantities on λ, ξ and m, when no confusion arises. For example, we will write $\rho_\lambda = \rho$, $\mathrm{RE}_\rho(k, \alpha) = \mathrm{RE}_{\rho_\lambda}(k, \alpha; \xi)$ or $C_\rho(A) = C^m_{\rho_\lambda}(A, \xi)$. The following lemma collects some simple consequences of these definitions.

Lemma B.5. *The following hold:*

(a) When ρ is nondecreasing,

- $A' \subset A \implies C_\rho(A') \subset C_\rho(A)$,
- $A' \subset A \implies \phi^2_\rho(Z, A') \geq \phi^2_\rho(Z, A)$,
- $Z \in \mathrm{RE}_\rho(k, \alpha) \implies \phi^2_\rho(Z, A) \geq \alpha^2,\ \forall A :\ |A| \leq k$,
- $k' \leq k \implies \mathrm{RE}_\rho(k, \alpha) \subset \mathrm{RE}_\rho(k', \alpha)$.

(b) $Z \in \mathrm{RE}_\rho(k, \alpha) \implies Z_S \in \mathrm{RE}_\rho(k \wedge |S|, \alpha)$.

The next result shows that we can control the generalized RE constants for Z_S uniformly by a suitable generalized RE constant for Z:

Lemma B.6. *If* $Z \in \mathrm{RE}_\rho(d, \alpha)$, *then*

$$\inf_{1 \leq j \leq m} \inf_{S \subset [m]_j} \inf_{\substack{A \subset S \\ |A| \leq d}} \phi^2_\rho(Z_S, A) \geq \alpha^2. \tag{B.10}$$

Proof. Fix j, $S \subset [m]_j$ and $A \subset S$ with $|A| \leq d$. By Lemma B.5(b), the assumption implies $Z_S \in \mathrm{RE}_\rho(d \wedge |S|, \alpha)$. Then, the last assertion in Lemma B.5(a) implies $Z_S \in \mathrm{RE}_\rho(|A|, \alpha)$, hence $\phi^2_\rho(Z_S, A) \geq \alpha^2$. Since the lower bound does not depend on A, S or j, we get the desired result. □

Next, we show that we can control RE constants for ρ by those for the ℓ_1 norm. Let us write $C^m_1(A, \xi)$ for the cone corresponding to $\rho = \|\cdot\|_1$, and similarly for the RE constants. We have the following lemma:

Lemma B.7. *Under Condition 2.2 on* ρ,

$$C^m_\rho(A, \xi) \subset \bigcup_{\substack{A' \subset [m] \\ |A'| = |A|}} C^m_1(A', \xi).$$

Proof of Lemma B.7. Fix nonzero $u \in C_\rho(A, \xi)$ and assume, without loss of generality that $|u_i| > 0$ for all $i \in [m]$; otherwise, we can inflate all the zero entries by $\epsilon > 0$, change ξ to $\xi + \rho(\epsilon)|A^c|/\rho(u_A)$, and let $\epsilon \to 0$ at the end.

Let $M = M(u)$ be the index set of the $|A|$ largest $|u_i|$. Then $\rho(u_{A^c}) \leq \xi\rho(u_A)$ implies $\rho(u_{M^c}) \leq \xi\rho(u_M)$ since ρ is nondecreasing. We note that $M = \{i :\ |u_i| \geq \tau\}$, for some $\tau > 0$, assuming $M^c \neq \emptyset$ without loss of generality. As a consequence of Condition 2.2, $x \mapsto \rho_\lambda(x)/x$ is nonincreasing. Then,

$$\tau\rho(|u_i|) \leq |u_i|\rho(\tau),\ i \in M,$$

with the reverse inequality for $i \in M^c$. Summing over M and M^c and rearranging, we have

$$\frac{\rho(u_M)}{\|u_M\|_1} \leq \frac{\rho(\tau)}{\tau} \leq \frac{\rho(u_{M^c})}{\|u_{M^c}\|_1},$$

or $\|u_{M^c}\|_1/\|u_M\| \leq \rho(u_{M^c})/\rho(u_M) \leq \xi$. Hence, $u \in C^m_1(M(u), \xi)$ where $|M(u)| = |A|$, which gives the desired result. □

As a consequence of Lemma B.7, $Z \in \mathrm{RE}_1(|A|, \alpha)$ implies $\phi^2_\rho(Z, A) \geq \alpha^2$, from which we get:

Lemma B.8. *Under Condition 2.2,* $\mathrm{RE}_1(d,\alpha) \subset \mathrm{RE}_\rho(d,\alpha)$.

In particular, the conclusion of Lemma B.6 holds under the (stronger) assumption $Z \in \mathrm{RE}_1(d,\alpha)$. In other words, to obtain uniform control over generalized restricted eigenvalues for all possible neighbourhood regression problems, it suffices to show that $\mathbf{X} \in \mathrm{RE}_1(d,\alpha)$ for some constant $\alpha > 0$. This is guaranteed by the following lemma, which is essentially a restatement of Corollary 1 in Raskutti et al. (2010):

Lemma B.9. *Assume* $\mathbf{X} \overset{iid}{\sim} \mathcal{N}_p(0,\Sigma)$ *for some* Σ *satisfying Condition 2.1. There exist universal constants* $c_0, c_1, c_2 > 0$*, such that if*

$$n > c_0 \frac{\sigma_{\max}^2 (1+\xi)^2}{r_{\min}(\Sigma)} d(\Sigma) \log p$$

then with probability at least $1 - c_1 \exp(-c_2 n)$*,*

$$\mathbf{X} \in \mathrm{RE}_1\big(d(\Sigma), \sqrt{r_{\min}(\Sigma)}; \xi\big).$$

Proposition B.4 now follows as a straightforward consequence of Lemmas B.6, B.8 and B.9.

Appendix C. Auxiliary results for sparsity

This section provides some additional results which are needed to prove Theorem 3.2 and Lemma 3.2, which are also proved in this section.

For any $\delta \in (0,1)$, $\lambda \geq 0$, $\delta_0 > 0$, and $\pi \in \mathbb{S}_p$, define the following event:

$$\mathcal{G}(\delta_0, \lambda; \pi) = \left\{ \frac{1}{2n} \|\mathbf{E}(\pi)\|_2^2 - \frac{1}{2n} \|\mathbf{E}(\widehat{\pi})\|_2^2 \leq \delta_0 \rho_\lambda(\widetilde{B}(\pi)) \right\}. \tag{C.1}$$

As in Appendix B, the idea will be to show that on this event (along with (B.1) and (B.2)), the desired conclusions hold. In Appendix C.5, we provide an explicit bound on the probability of $\mathcal{G}(\delta_0, \lambda; \pi)$.

C.1. Some intermediate lemmas.

Recall the definitions of $\mathcal{E}(\delta,\lambda)$ and $\mathcal{R}(\delta)$ in (B.1)–(B.2). We start with the following extension of GW bounds:

Lemma C.1. *Let* $\widehat{\Delta} := \widehat{B} - \widetilde{B}(\widehat{\pi})$*. On* $\mathcal{E}(\delta,\lambda)$*, we have*

$$\frac{1}{n} \left| \mathrm{tr}\left(\mathbf{E}(\widehat{\pi})^T \mathbf{X} \widehat{\Delta} \right) \right| < \delta \Big[\frac{1}{2n} \|\mathbf{X}\widehat{\Delta}\|_2^2 + \rho_\lambda(\widehat{\Delta}) \Big]. \tag{C.2}$$

Proof. Let $\widehat{\Delta}_j := \widehat{\beta}_j - \widetilde{\beta}_j(\widehat{\pi})$ be the jth column of $\widehat{\Delta}$. Then

$$\frac{1}{n} \left| \mathrm{tr}\left(\mathbf{E}(\widehat{\pi})^T \mathbf{X} \widehat{\Delta} \right) \right| \leq \frac{1}{n} \sum_{j=1}^p |\langle \widetilde{\boldsymbol{\varepsilon}}_j(\widehat{\pi}), \mathbf{X}\widehat{\Delta}_j \rangle|. \tag{C.3}$$

According to (B.1), on $\mathcal{E}(\delta,\lambda)$, we have $(\widetilde{\varepsilon}_j(S), \mathbf{X}_S) \in \mathrm{GW}^{\circ}_{\rho_\lambda}(\delta)$ for all $S \subset [p]_j$. In particular, applying with $S = S_j(\widehat{\pi})$ and using $u = \widehat{\Delta}_j$ in the Definition 5.3 of GW, we have

$$\frac{1}{n}|\langle \widetilde{\varepsilon}_j(\widehat{\pi}), \mathbf{X}\widehat{\Delta}_j\rangle| < \delta\Big[\frac{1}{2n}\|\mathbf{X}\widehat{\Delta}_j\|_2^2 + \rho_\lambda(\widehat{\Delta}_j)\Big], \quad \forall j$$

Summing over j and plugging into (C.3) yields (C.2). □

For any matrix $A = (a_{ij}) \in \mathbb{R}^{p\times p}$ and $S \subset [p]\times[p]$, let $A_{\langle S\rangle}$ denote the $p \times p$ matrix formed by zero-ing the elements outside of S, i.e.

$$(A_{\langle S\rangle})_{ij} = \begin{cases} a_{ij}, & (i,j) \in S, \\ 0, & (i,j) \notin S. \end{cases}$$

In analogy with Condition 3.1 on signal strength, let us define

$$\tau_\lambda(\alpha;\Sigma) := \inf\Big\{\tau : \frac{\rho'_\lambda(0+)^2}{\rho_\lambda(\tau)} \le \frac{r_{\min}(\Sigma)}{\alpha}\Big\} \tag{C.4}$$

where we often suppress the dependence on Σ. Note that we can write Condition 3.1 equivalently as $\tau_* \ge \tau_\lambda(a_1)$.

Lemma C.2. *Assume that ρ_λ satisfies Condition 2.2 and*

$$\tau_* \ge \tau_\lambda\Big(\frac{2\xi}{1-\delta_1}\Big), \quad \textit{for some } \delta_1 \in (0,1) \tag{C.5}$$

where $\xi = \xi(\delta)$ is defined by (B.4). Then, on $\mathcal{R}(\delta)\cap\mathcal{E}(\delta,\lambda)$,

$$\rho_\lambda(\widehat{B}) \;\ge\; \delta_1\rho_\lambda(\widetilde{B}(\widehat{\pi})) + \rho_\lambda\Big((\widehat{B}-\widetilde{B}(\widehat{\pi}))_{\langle \mathrm{supp}(\widetilde{B}(\widehat{\pi}))^c\rangle}\Big). \tag{C.6}$$

Proof. To lighten the notation, let $\Delta = \widehat{B} - \widetilde{B}(\widehat{\pi})$, $S_1 = \mathrm{supp}(\widetilde{B}(\widehat{\pi}))$, $\Delta_1 = \Delta_{\langle S_1\rangle}$, and $\Delta_2 = \Delta_{\langle S_1^c\rangle}$. We have

$$\rho_\lambda(\Delta_1) \le \rho'_\lambda(0+)\|\Delta_1\|_1 \le \rho'_\lambda(0+)\|\widetilde{B}(\widehat{\pi})\|_0^{1/2}\|\Delta_1\|_2. \tag{C.7}$$

Since we are on $\mathcal{R}(\delta)\cap\mathcal{E}(\delta,\lambda)$, Proposition B.1 yields the ℓ_2 deviation bound (B.5), which we use with $S = S_j(\widehat{\pi})$. Plugging into (C.7) and using $\|\Delta_1\|_2 \le \|\Delta\|_2$,

$$\rho_\lambda(\Delta_1) \le \big[\rho'_\lambda(0+)\big]^2\frac{2\xi}{r_{\min}(\Sigma)}\|\widetilde{B}(\widehat{\pi})\|_0. \tag{C.8}$$

Trivially, we have $\rho_\lambda(\widetilde{B}(\widehat{\pi})) \ge \rho_\lambda(\tau_*)\|\widetilde{B}(\widehat{\pi})\|_0$, so that by (C.8)

$$\rho_\lambda(\Delta_1) \le \Big[\frac{\rho'_\lambda(0+)^2}{\rho_\lambda(\tau_*)}\frac{2\xi}{r_{\min}(\Sigma)}\Big]\rho_\lambda(\widetilde{B}(\widehat{\pi})) \;\le\; (1-\delta_1)\rho_\lambda(\widetilde{B}(\widehat{\pi})), \tag{C.9}$$

where the last inequality follows from (C.5). Finally, note that

$$\begin{aligned}\rho_\lambda(\widehat{B}) \;&\ge\; \rho_\lambda(\widetilde{B}(\widehat{\pi})) + \rho_\lambda(\Delta_2) - \rho_\lambda(\Delta_1) \\ &\ge\; \delta_1\rho_\lambda(\widetilde{B}(\widehat{\pi})) + \rho_\lambda(\Delta_2).\end{aligned}$$

where the first inequality is by arguments similar to those leading to (A.4) and the second is by (C.9). □

The conclusion of Lemma C.2 is stronger than what we need in the sequel. We only use the weaker inequality $\rho_\lambda(\widehat{B}) \geq \delta_1 \rho_\lambda(\widetilde{B}(\widehat{\pi}))$ implied by (C.6).

C.2. **A sparsity bound.** Proposition C.3 below is the main ingredient used in the proof of Theorem 3.2. This result provides an explicit, nonasymptotic relationship between the "weak" sparsity of $\{\widehat{B}, \widetilde{B}(\widehat{\pi}), \widetilde{B}(\pi_0)\}$ as measured by the regularizer ρ_λ.

Proposition C.3. *Assume $n > 8\,(d+1)\log p$. Under Condition 2.2 on ρ_λ, further assume*

$$\tau_* \geq \tau_\lambda\Big(\frac{2(1+\delta)}{1-3\delta}\Big) \quad \textit{for some } \delta \in (0, 1/3).$$

Then, given π_0 satisfying Condition 3.2 for some $a_2 > 0$, we have

$$\frac{2\delta}{1-\delta}\rho_\lambda(\widetilde{B}(\widehat{\pi})) \overset{(i)}{\leq} \rho_\lambda(\widehat{B}) \overset{(ii)}{\leq} \frac{2}{1-\delta}\Big(1+\frac{10}{a_2}\Big)\rho_\lambda(\widetilde{B}(\pi_0)), \tag{C.10}$$

with probability at least $1 - c_1 e^{-c_2 \min\{n,\,(d+1)\log p\}} - p\binom{p}{d}\mathbb{E}e^{-\psi_\lambda(\mathbf{X},\sigma^2_{\max};\delta)}$.

Proof. Recall the definition of $\mathcal{G}(\delta_0, \lambda; \pi)$ in (C.1). Fix some π_0 satisfying Condition 3.2 with $a_2 > 0$. Taking (arbitrarily) $C = 1$ and $\delta_0 = 10/a_2$ in Proposition C.4, we have

$$\mathbb{P}\left[\mathcal{G}(\delta_0, \lambda; \pi_0)^c\right] \leq 2e^{-(d+1)\log p}.$$

Combined with Propositions B.4 and B.3, we obtain

$$\begin{aligned}&\mathbb{P}\left(\mathcal{G}(\delta_0, \lambda; \pi_0) \cap \mathcal{E}(\delta, \lambda) \cap \mathcal{R}(\delta)\right)\\ &\quad\geq 1 - c_1 \exp(-c_2 \min\{n, (d+1)\log p\}) - p\binom{p}{d}\mathbb{E}\exp(-\psi_\lambda(\mathbf{X}, \sigma^2_{\max}; \delta)).\end{aligned}$$

Thus, we may assume we are on $\mathcal{G}(\delta_0, \lambda; \pi_0) \cap \mathcal{E}(\delta, \lambda) \cap \mathcal{R}(\delta)$. Since we are on $\mathcal{E}(\delta, \lambda)$, we can combine Lemma C.1 with Lemma 5.7 (applied with $\pi = \pi_0$) to deduce (recall $\widehat{\Delta} := \widehat{B} - \widetilde{B}(\widehat{\pi})$)

$$\begin{aligned}\frac{1}{2n}\|\mathbf{X}\widehat{\Delta}\|_2^2 + \rho_\lambda(\widehat{B}) \;&\leq \frac{\delta}{2n}\|\mathbf{X}\widehat{\Delta}\|_2^2 + \delta\rho_\lambda(\widehat{\Delta})\\ &\quad + \frac{1}{2n}\|\mathbf{E}(\pi_0)\|_2^2 - \frac{1}{2n}\|\mathbf{E}(\widehat{\pi})\|_2^2 + \rho_\lambda(\widetilde{B}(\pi_0)).\end{aligned}$$

Dropping the prediction loss terms (those involving $\|\mathbf{X}\widehat{\Delta}\|_2^2$), and using that we are on $\mathcal{G}(\delta_0, \lambda; \pi_0)$ to bound $\frac{1}{2n}\|\mathbf{E}(\pi_0)\|_2^2 - \frac{1}{2n}\|\mathbf{E}(\widehat{\pi})\|_2^2$, we have after rearranging,

$$\begin{aligned}\rho_\lambda(\widehat{B}) &\leq (1+\delta_0)\rho_\lambda(\widetilde{B}(\pi_0)) + \delta\rho_\lambda(\widehat{B} - \widetilde{B}(\widehat{\pi}))\\ &\leq (1+\delta_0)\rho_\lambda(\widetilde{B}(\pi_0)) + \delta\rho_\lambda(\widetilde{B}(\widehat{\pi})) + \delta\rho_\lambda(\widehat{B}).\end{aligned} \tag{C.11}$$

Let $\delta_1 = 2\delta/(1-\delta)$, so that $\xi/(1-\delta_1) = (1+\delta)/(1-3\delta)$ (cf. (B.4)). Furthermore, since $\delta < 1/3$ by assumption, $\delta_1 < 1$, so that Lemma C.2 implies $\rho_\lambda(\widehat{B}) \geq \delta_1 \rho_\lambda(\widetilde{B}(\widehat{\pi}))$ which gives (i) in (C.10).

Since $\rho_\lambda(\widetilde{B}(\widehat{\pi})) \le (1/\delta_1)\rho_\lambda(\widehat{B})$, the bounds in (C.11) imply that

$$\rho_\lambda(\widehat{B}) \le (1+\delta_0)\rho_\lambda(\widetilde{B}(\pi_0)) + \frac{\delta}{\delta_1}\rho_\lambda(\widehat{B}) + \delta\rho_\lambda(\widehat{B}).$$

Rearranging we get

$$\rho_\lambda(\widehat{B}) \le \big[1-\delta(1+\delta_1)/\delta_1\big]^{-1}(1+\delta_0)\rho_\lambda(\widetilde{B}(\pi_0)).$$

We have $[1-\delta(1+\delta_1)/\delta_1]^{-1}(1+\delta_0) = \frac{2}{1-\delta}(1+\frac{10}{a_2})$, using $\delta_0 = 10/a_2$ and $\delta_1 = 2\delta/(1-\delta)$ as before. This proves (ii) in (C.10). □

C.3. **Proof of Theorem 3.2.** For regularizers satisfying Conditions 2.2 and 2.3, the desired bound follows by taking $\delta = (a_1-2)/(3a_1+2) \in (0,1/3)$ in Proposition C.3, and using Proposition B.2 to complete the probability bound.

To deduce the ℓ_0 bound, consider the case where ρ_λ is also ℓ_0-compatible. Condition 3.1 implies

$$\rho_\lambda(\widetilde{B}(\widehat{\pi})) \ge a_1\frac{\rho'_\lambda(0+)^2}{r_{\min}(\Sigma)}\|\widetilde{B}(\widehat{\pi})\|_0,$$

while on the other hand, ℓ_0-compatibility (Definition 2.4) gives $\rho_\lambda(\widetilde{B}(\pi_0)) \le \overline{\rho}_0\lambda^2\|\widetilde{B}(\pi_0)\|_0$. Combining these with (C.10) yields

$$\begin{aligned} a_1\frac{\rho'_\lambda(0+)^2}{r_{\min}(\Sigma)}\|\widetilde{B}(\widehat{\pi})\|_0 &\le \frac{1}{\delta}\Big(1+\frac{10}{a_2}\Big)\overline{\rho}_0\,\lambda^2\|\widetilde{B}(\pi_0)\|_0 \\ \implies \|\widetilde{B}(\widehat{\pi})\|_0 &\le \kappa_3\cdot\Big[\frac{\lambda^2}{\rho'_\lambda(0+)^2}\Big]\|\widetilde{B}(\pi_0)\|_0, \end{aligned}$$

as desired. Here, $\kappa_3 = \kappa_3(\Sigma) = \frac{1}{a_1}\big(\frac{3a_1+2}{a_1-2}\big)\big(1+\frac{10}{a_2}\big)\overline{\rho}_0 r_{\min}(\Sigma)$, using our earlier choice of δ.

C.4. **Proof of Lemma 3.2.** Assume that we are on the event $\mathcal{R}(\delta)\cap\mathcal{E}(\delta,\lambda)$. As in the proof of Theorem 3.2, choose $\delta = (a_1-2)/(3a_1+2) \in (0,1/3)$. Applying (C.2) to Lemma 5.7 yields

$$\frac{1}{2n}\|\mathbf{E}(\widehat{\pi})\|_2^2 + \rho_\lambda(\widehat{B}) \le \frac{1}{2n}\|\mathbf{E}(\pi)\|_2^2 + \rho_\lambda(\widetilde{B}(\pi)) + \delta\rho_\lambda(\widehat{B}-\widetilde{B}(\widehat{\pi})) \tag{C.12}$$

for any $\pi \in \mathbb{S}_p$. Lemma C.1 combined with and Propositions B.2 and B.3 imply that (C.12) holds with probability at least $1-c_1\exp(-c_2(d+1)\log p)$.

Invoking the subadditivity of ρ_λ and re-arranging,

$$\frac{1}{2n}\|\mathbf{E}(\widehat{\pi})\|_2^2 + (1-\delta)\rho_\lambda(\widehat{B}) \le \frac{1}{2n}\|\mathbf{E}(\pi)\|_2^2 + \rho_\lambda(\widetilde{B}(\pi)) + \delta\rho_\lambda(\widetilde{B}(\widehat{\pi})). \tag{C.13}$$

Now apply Lemma C.2 with $\delta_1 = 2\delta/(1-\delta)$ to deduce that

$$\frac{1}{2n}\|\mathbf{E}(\widehat{\pi})\|_2^2 + \delta\rho_\lambda(\widetilde{B}(\widehat{\pi})) \le \frac{1}{2n}\|\mathbf{E}(\pi)\|_2^2 + \rho_\lambda(\widetilde{B}(\pi)). \tag{C.14}$$

Taking $u = \sqrt{4(d+1)\log p}$ in Lemma C.6 and observing that $1 - h_n(u) > 1/2 > \delta$ (cf. (C.16)) for $n > 5(d+1)\log p$, we have

$$\begin{aligned}\frac{1}{2}\operatorname{tr}\widetilde{\Omega}(\widehat{\pi}) + \rho_\lambda(\widetilde{B}(\widehat{\pi})) &\le \frac{1+H_n(u)}{\delta}\cdot\left[\frac{1}{2}\operatorname{tr}\widetilde{\Omega}(\pi) + \rho_\lambda(\widetilde{B}(\pi))\right] \\ &\le \frac{1}{\delta}\left(1 + 6\sqrt{\frac{(d+1)\log p}{n}}\right)\cdot\left[\frac{1}{2}\operatorname{tr}\widetilde{\Omega}(\pi) + \rho_\lambda(\widetilde{B}(\pi))\right].\end{aligned}$$

Combined with the probability bound implied by Propositions B.2 and B.3, this holds with probability at least $1 - c_1\exp(-c_2(d+1)\log p)$ (c_1, c_2 possibly different from above). Since π was arbitrary, substituting $\delta = (a_1 - 2)/(3a_1 + 2)$ completes the proof.

C.5. **A bound on the sample residuals.** In this section, we prove the following result, which is used in the proof of Proposition C.3:

Proposition C.4. *Assume $n > 4(C+1)(d+1)\log p$ for some $C > 0$ and let π_0 be a minimum-trace permutation such that*

$$\frac{\rho_\lambda(\widetilde{B}(\pi_0))}{\operatorname{tr}\widetilde{\Omega}(\pi_0)} \ge \frac{1}{\delta_0}\sqrt{\frac{50(C+1)(d+1)\log p}{n}}. \tag{C.15}$$

Then for any $\delta_0 > 0$, $\mathbb{P}(\mathcal{G}(\delta_0, \lambda; \pi_0)) \ge 1 - 2e^{-C(d+1)\log p}$, i.e.

$$\mathbb{P}\left(\frac{1}{2n}\|\mathbf{E}(\pi_0)\|_2^2 - \frac{1}{2n}\|\mathbf{E}(\widehat{\pi})\|_2^2 > \delta_0\rho_\lambda(\widetilde{B}(\pi_0))\right) \le 2e^{-C(d+1)\log p}.$$

Define two functions by

$$h_n(u) := -\frac{u^2}{n} + \frac{2u}{\sqrt{n+1}} + \frac{1}{n+1}, \quad H_n(u) := \frac{u^2}{n} + \frac{2u}{\sqrt{n}}. \tag{C.16}$$

These functions bound the deviations in the normed residuals $\widetilde{\varepsilon}_j(\pi)$, and will be used repeatedly in the sequel. We note that

$$H_n(u) + h_n(u) \le \frac{5u}{\sqrt{n}}, \quad u \ge n^{-1/2}. \tag{C.17}$$

Lemma C.5. *Suppose $\mathbf{w} \sim \mathcal{N}_n(0, \sigma^2 I_n)$. Then for any $0 < u < n/\sqrt{n+1}$,*

$$\sigma^2\Big(1 - h_n(u)\Big) \le \frac{1}{n}\|\mathbf{w}\|_2^2 \le \sigma^2\Big(1 + H_n(u)\Big) \tag{C.18}$$

with probability at least $1 - 2e^{-u^2/2}$.

Proof. For $z \sim \mathcal{N}_n(0, I_n)$, we have the following useful bounds (see, e.g., Gordon, 1988, Corollary 1.2):

$$\frac{n}{\sqrt{n+1}} \le \mathbb{E}\|z\|_2 = \sqrt{2}\frac{\Gamma(\frac{n+1}{2})}{\Gamma(\frac{n}{2})} \le \sqrt{n}.$$

Gaussian concentration implies that for any $u > 0$, both

$$\big\{ \|\mathbf{w}\|_2 \le \sigma(n/\sqrt{n+1} - u)\big\}, \quad \text{and} \quad \big\{ \|\mathbf{w}\|_2 \ge \sigma(\sqrt{n} + u)\big\}$$

hold with probability at most $e^{-u^2/2}$. Thus,

$$\mathbb{P}\left(\sigma^2\Big(\frac{n}{\sqrt{n+1}} - u\Big)^2 \le \|\mathbf{w}\|_2^2 \le \sigma^2\big(\sqrt{n} + u\big)^2\right) \ge 1 - 2e^{-u^2/2}. \tag{C.19}$$

Re-writing (C.19) using (C.16) yields the desired result. □

Lemma C.6. *Suppose* $\mathbf{X} \overset{iid}{\sim} \mathcal{N}_p(0, \Sigma)$. *Then for any* $\pi \in \mathbb{S}_p$ *and* $0 < u < n/\sqrt{n+1}$,

$$\frac{1}{2}\operatorname{tr}\widetilde{\Omega}(\pi)\Big(1 - h_n(u)\Big) \le \frac{1}{2n}\|\mathbf{E}(\pi)\|_2^2 \le \frac{1}{2}\operatorname{tr}\widetilde{\Omega}(\pi)\Big(1 + H_n(u)\Big) \tag{C.20}$$

with probability at least $1 - 2p\binom{p}{d}e^{-u^2/2}$.

Proof. Note that for any $\pi \in \mathbb{S}_p$,

$$\frac{1}{2n}\|\mathbf{E}(\pi)\|_2^2 = \frac{1}{2n}\sum_{j=1}^{p}\|\widetilde{\boldsymbol{\varepsilon}}_j(\pi)\|_2^2 = \frac{1}{2n}\sum_{j=1}^{p}\|\widetilde{\boldsymbol{\varepsilon}}_j(S_j(\pi))\|_2^2\,. \tag{C.21}$$

Thus it suffices to bound the deviations in $\|\widetilde{\boldsymbol{\varepsilon}}_j(S)\|_2$ for $S \subset [p]_j$. Consider the following events

$$\mathcal{G}_j(S) := \left\{\frac{\omega_j^2(S)}{2}\Big(1 - h_n(u)\Big) \le \frac{1}{2n}\|\widetilde{\boldsymbol{\varepsilon}}_j(S)\|_2^2 \le \frac{\omega_j^2(S)}{2}\Big(1 + H_n(u)\Big)\right\}$$

and let $\mathcal{G} := \bigcap_{j=1}^{p}\bigcap_{S\subset[p]_j}\mathcal{G}_j(S)$. By Lemma C.5, we have $\mathbb{P}(\mathcal{G}_j(S)) \ge 1 - 2e^{-u^2/2}$, for all $S \in [p]_j$. By a monotonicity argument (cf. (5.3)), we have $\mathcal{G} = \bigcap_{j=1}^{p}\bigcap_{S\in m_j(\Sigma)}\mathcal{G}_j(M_j(S))$. Applying the union bound and using (2.4),

$$\mathbb{P}(\mathcal{G}^c) \le 2p\binom{p}{d}e^{-u^2/2}. \tag{C.22}$$

Summing the inequalities defining $\mathcal{G}_j(S_j(\pi))$, over j, we conclude that (C.20) holds on $\mathcal{G}$. The proof is complete. □

Consider the (random) collection of permutations

$$\mathbb{S}_p^0 = \mathbb{S}_p^0(\delta_0; u) := \Big\{\pi \in \mathbb{S}_p : \frac{1}{2}\operatorname{tr}\widetilde{\Omega}(\pi)\Big[1 + H_n(u)\Big] - \frac{1}{2}\operatorname{tr}\widetilde{\Omega}(\widehat{\pi})\Big[1 - h_n(u)\Big] \le \delta_0\rho_\lambda(\widetilde{B}(\pi))\Big\}.$$

Lemma C.7. *For any* $\pi \in \mathbb{S}_p^0(\delta_0; u)$ *and* $0 < u < n/\sqrt{n+1}$, *we have*

$$\mathbb{P}\left(\frac{1}{2n}\|\mathbf{E}(\pi)\|_2^2 - \frac{1}{2n}\|\mathbf{E}(\widehat{\pi})\|_2^2 > \delta_0\rho_\lambda(\widetilde{B}(\pi))\right) \le 2p\binom{p}{d}e^{-u^2/2}.$$

Proof. Lemma C.6 implies that

$$\frac{1}{2n}\|\mathbf{E}(\pi)\|_2^2 - \frac{1}{2n}\|\mathbf{E}(\widehat{\pi})\|_2^2 \leq \frac{1}{2}\operatorname{tr}\widetilde{\Omega}(\pi)\Big[1 + H_n(u)\Big] - \frac{1}{2}\operatorname{tr}\widetilde{\Omega}(\widehat{\pi})\Big[1 - h_n(u)\Big]$$

with probability at least $1 - 2p\binom{p}{d}e^{-u^2/2}$. Since $\pi \in \mathbb{S}_p^0$, the right-side is bounded above by $\delta_0\rho_\lambda\widetilde{B}(\pi)$ by definition, which establishes the claim. □

Proof of Proposition C.4. Lemma C.7 implies that for a choice of $u = \sqrt{2(C+1)(d+1)\log p}$, we have

$$\begin{aligned}\mathbb{P}\left(\frac{1}{2n}\|\mathbf{E}(\pi)\|_2^2 - \frac{1}{2n}\|\mathbf{E}(\widehat{\pi})\|_2^2 > \delta_0\rho_\lambda(\widetilde{B}(\pi))\right) &\leq 2p\binom{p}{d}e^{-(C+1)(d+1)\log p} \\ &\leq 2e^{-C(d+1)\log p}\end{aligned}$$

for any $\pi \in \mathbb{S}_p^0$. Thus the claim will follow if we can show that $\pi_0 \in \mathbb{S}_p^0$.

Since $\operatorname{tr}\widetilde{\Omega}(\pi_0) \leq \operatorname{tr}\widetilde{\Omega}(\widehat{\pi})$, we have

$$\begin{aligned}\operatorname{tr}\widetilde{\Omega}(\pi_0)\Big[1 + H_n(u)\Big] - \operatorname{tr}\widetilde{\Omega}(\widehat{\pi})\Big[1 - h_n(u)\Big] &\leq \operatorname{tr}\widetilde{\Omega}(\pi_0)\Big[H_n(u) + h_n(u)\Big] \\ &\overset{(i)}{\leq} \operatorname{tr}\widetilde{\Omega}(\pi_0)\sqrt{\frac{50(C+1)(d+1)\log p}{n}} \\ &\overset{(ii)}{\leq} \delta_0\rho_\lambda(\widetilde{B}(\pi_0))\end{aligned}$$

where (i) follows by using (C.17) with $u = \sqrt{2(C+1)(d+1)\log p}$, and (ii) follows from assumption (C.15). Hence, $\pi_0 \in \mathbb{S}_p^0$ and the proof is complete. □

References

B. Aragam and Q. Zhou. Concave penalized estimation of sparse Gaussian Bayesian networks. *The Journal of Machine Learning Research*, 16:1–56, 2015.

P. J. Bickel, Y. Ritov, and A. B. Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *The Annals of Statistics*, 37(4):1705–1732, 2009.

P. Bühlmann, J. Peters, J. Ernest, et al. CAM: Causal additive models, high-dimensional order search and penalized regression. *The Annals of Statistics*, 42(6):2526–2556, 2014.

J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96 (456):1348–1360, 2001.

F. Fu and Q. Zhou. Learning sparse causal Gaussian networks with experimental intervention: Regularization and coordinate descent. *Journal of the American Statistical Association*, 108(501):288–300, 2013.

F. Fu, J. Gu, and Q. Zhou. Adaptive penalized estimation of directed acyclic graphs from categorical data. *arXiv preprint arXiv:1403.2310*, 2014.

Y. Gordon. On Milman's inequality and random subspaces which escape through a mesh in $\mathbb{R}^n$. *Geometric Aspects of Functional Analysis*, pages 84–106, 1988.

J. Huang, P. Breheny, and S. Ma. A selective review of group selection in high-dimensional models. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 27(4), 2012.

P.-L. Loh and P. Bühlmann. High-dimensional learning of linear causal networks via inverse covariance estimation. *Journal of Machine Learning Research*, 15:3065–3105, 2014.

P.-L. Loh and M. J. Wainwright. Regularized m-estimators with nonconvexity: Statistical and algorithmic theory for local optima. *arXiv preprint arXiv:1305.2436*, 2013.

N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the Lasso. *The Annals of Statistics*, 34(3):1436–1462, 2006.

J. Peters and P. Bühlmann. Identifiability of gaussian structural equation models with equal error variances. *Biometrika*, page ast043, 2013.

G. Raskutti, M. J. Wainwright, and B. Yu. Restricted eigenvalue properties for correlated Gaussian designs. *The Journal of Machine Learning Research*, 11:2241–2259, 2010.

G. Raskutti, M. J. Wainwright, and B. Yu. Minimax rates of estimation for high-dimensional linear regression over ℓ_q-balls. *Information Theory, IEEE Transactions on*, 57(10):6976–6994, 2011.

P. Ravikumar, M. J. Wainwright, J. D. Lafferty, et al. High-dimensional ising model selection using ℓ_1-regularized logistic regression. *The Annals of Statistics*, 38(3):1287–1319, 2010.

M. Schmidt, A. Niculescu-Mizil, and K. Murphy. Learning graphical model structure using L1-regularization paths. In *AAAI*, volume 7, pages 1278–1283, 2007.

A. Shojaie and G. Michailidis. Penalized likelihood methods for estimation of sparse high-dimensional directed acyclic graphs. *Biometrika*, 97(3): 519–538, 2010.

S. van de Geer and P. Bühlmann. ℓ_0-penalized maximum likelihood for sparse directed acyclic graphs. *The Annals of Statistics*, 41(2):536–567, 2013.

C.-H. Zhang. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38(2):894–942, 2010.

C.-H. Zhang and J. Huang. The sparsity and bias of the Lasso selection in high-dimensional linear regression. *The Annals of Statistics*, pages 1567–1594, 2008.

C.-H. Zhang and T. Zhang. A general theory of concave regularization for high-dimensional sparse estimation problems. *Statistical Science*, 27(4): 576–593, 2012.